Trust Me: Communication and Competition in a Psychological Game*

Marina Agranov[†] Utteeyo Dasgupta[‡] Andrew Schotter[§]

December 13, 2023

Abstract

We study, both theoretically and experimentally, a communication game with and without seller competition and embed it in a psychological-game framework where players experience costs for lying, misleading others, and being disappointed. We derive the equilibrium predictions of this model, compare them to the setting without psychological payoffs, and test these predictions in a laboratory experiment, in which we induce both material and psychological payoffs. We find that the setting in which players have both material and psychological payoffs features more trade, trades goods of marginally better quality, and does so without welfare losses to either side of the market relative to the setting with material payoffs only. However, the introduction of competition counteracts this improvement and lowers welfare for both sides of the market. This happens due to a surge in dishonesty by sellers in the competitive setting and the buyers' inability to detect this deception.

1 Introduction

Since the seminal paper by Crawford and Sobel (1982), economists have devoted considerable attention to communication games. These games typically involve an informed sender who sends a message to a less informed receiver, who then takes an action that determines the payoffs to both parties. The question investigated in these models is the informativeness of the equilibrium messages sent by senders as a function of the divergence of their preferences from receivers over material outcomes. In these games, senders may lie and deceive others (Kartik, 2009; Sobel, 2020) but in the models underlying these games, there is no room for feeling guilty when one misleads others or feeling disappointed when one is lied to. This raises the question of whether the equilibria of these games would be more informative (honest) if senders suffered from both lying and deception aversion (guilt) and receivers could feel disappointment.

*We thank the seminar participants at Princeton and Navin Kartik for helpful comments and suggestions.

†Division of the Humanities and Social Sciences, California Institute of Technology and NBER. Email: magranov@hss.caltech.edu.

‡Department of Economics, Fordham University. Email: utteeyodasgupta@gmail.com.

§Department of Economics, New York University. Email: andrew.schotter@nyu.edu.

Our paper aims to address this question by modeling a market as a sender-receiver game with psychological costs and testing it experimentally in the lab.[1] The inclusion of these costs can help regulate these markets, diminish deception, and establish equilibria that are more informative than the pooling equilibrium that would likely prevail without them. Furthermore, we investigate how these markets respond to the introduction of competition among senders, with the hope that such competition may further enhance information transmission between the senders and the receiver.

Setup. Our experiment tests three games against their equilibrium predictions. The first game is the standard sender-receiver (seller-buyer) game without competition and without any psychological payoffs. It involves two players: the seller and the buyer. The seller owns one unit of a product that is either of high or low quality and wants to sell it to the buyer. The buyer is interested in purchasing a high-quality product and not a low-quality one. The situation is complicated by the fact that the buyer cannot distinguish between the high- and the low-quality product without actually purchasing it, and instead has to rely on the messages sent by the seller. In this game, players have only material payoffs.

In a second game, in addition to the material payoffs, players receive psychological payoffs, which depend on their types. Although psychological payoffs can be multi-dimensional, we focus on the most prominent ones identified in the literature. Sellers can suffer from lying when they misrepresent the quality of the goods they own, and can also feel guilty if they mislead the buyers about the quality of their goods. Buyers can suffer from disappointment when their beliefs about the quality of the goods are misplaced.

Our third game — a game with competition — is identical to our second game except that we add a second seller who competes with the first one. The competition happens via communication, where each seller sends his own message to the buyer. The buyer picks one of the sellers based on the received messages and decides whether to purchase the product from the selected seller or not.

Theoretical Predictions. The introduction of psychological payoffs to a communication game without competition is potentially beneficial in terms of generating positive trade in equilibrium. Indeed, the absence of trade is the only possible equilibrium outcome in the first game with only material payoffs. Conversely, in the second game, in which players have both material and psychological payoffs, in addition to the pooling equilibrium, multiple informative equilibria exist. In these equilibria, sellers' messages are partially informative, and, as a result, trade occurs with positive probability. For the parameters utilized in our experiment, the second game features two informative equilibria, each resulting in higher expected payoffs for both the seller and the buyer compared to the pooling equilibrium with no trade. Furthermore, these two informative equilibria can be ranked based on the amount of information transmitted during the communication stage, with higher message informativeness corresponding to greater trade frequency, higher quality of sold goods, and a higher payoff for the buyer.

The effect of introducing competition into a setting with both material and psychological payoffs is theoretically less clear. The third game with competition features the same set of equilibria as the second game without competition, i.e., the pooling equilibrium with no trade and the two informative equilibria with different levels of message informativeness. The theory is, however, silent regarding which equilibrium is more likely to be played. If the same or a more informative equilibrium is played in the presence of competition, buyers stand to benefit from it. However, if

[1] Our theoretical model belongs to the class of psychological games, in which players' payoffs depend not only on their actions but also on their beliefs.

sellers with low-quality products feel compelled to lie more due to competition, this will diminish messages' informativeness and negatively affect the buyers' ability to select a better seller to engage in trade. The latter effect can ultimately lead to a selection of a less informative equilibrium in which buyers' welfare is lower. Which result holds is ultimately an empirical question. Consequently, we turn to controlled laboratory experiments to help us sort things out by comparing the results in our three games performed in the lab.

It is worth noting that lying aversion, deriving disutility from not telling the truth, is not sufficient to establish all of the equilibria we sustain in our model. This is true because lying, per se, is not a psychological-game force, since the disutility from lying does not depend on players' beliefs about others and simply captures the fact that sellers dislike being dishonest, i.e., sending a message that does not match the quality of a product they own.[2] Absent psychological payoffs, which link players' payoffs to their beliefs, both games with and without competition admit at most *one* informative equilibrium, which makes our setting considerably less interesting.[3]

Main Features of the Experimental Design. In terms of experimental design, our paper is unique in two ways. First, to test the equilibrium predictions described above, we induce the psychological costs assumed by the theory, rather than estimating them or inferring them from the data. To do so, we utilize the classical experimental approach of induced value pioneered by Smith (1976) and, in accordance with our theoretical model, impose monetary costs on sellers when they lie and disappoint the buyer, and on the buyer when she is misled by the seller. The goal of the experiment is to observe how such payoffs affect game outcomes in the presence and absence of competition between sellers.[4]

Second, in order to give subjects experience with the game, we ask them to choose a strategy once in each of the ten rounds of the experiment and then simulate the outcomes of the experiment given these strategy choices ten times. This approach allows our subjects to receive feedback from 100 rounds throughout the experiment while making only ten decisions, offering them a streamlined and effective experience in the game.

Experimental Results. Our experimental findings support some predictions of the theory and refute others. In terms of market performance, introducing psychological payoffs in markets with no competition increases trade and marginally improves the quality of purchased goods.[5] When competition is introduced into markets with psychological payoffs, trade further increases but at the cost of reduced product quality.

Welfare comparisons across markets reflect differences in market performance but also account for psychological costs induced in the second and third games. Our data shows that buyers' welfare is similar in markets with material payoffs alone and those with additional psychological payoffs but no competition. The sellers' welfare is also comparable in these two types of markets. At the

[2]The recent experimental literature has convincingly documented that people possess an intrinsic aversion to lying when messages are cheap talk (Gneezy, 2005; Hurkens and Kartik, 2009; Sanchez-Pages and Vorsatz, 2007) and that people are reluctant to tell even white lies, which benefit both the person telling the lie and the one to whom the lie is told (Erat and Gneezy, 2012).

[3]We discuss in more detail what happens in the model without guilt and disappointment aversion in Section 2.5.

[4]In Section 3.2, we discuss in detail the challenges associated with using the induced value method to induce psychological payoffs, how we overcome these challenges, and the extent to which the experimenter can successfully control subjects' home-grown psychological costs.

[5]We show in Section 4 that the average quality of sold goods in the second game is higher than that in the first game but not significantly so.

same time, both buyers and sellers earn significantly lower payoffs in markets with competition and psychological payoffs compared to the markets with neither. Finally, the introduction of competition into markets with psychological payoffs leads to lower welfare for both buyers and sellers.

Hence, the punchline of our paper is simple: Settings where people experience psychological costs, feel morally accountable for their actions but face no competition, feature more trade, trade goods of marginally better quality, and do so without welfare losses to either side of the market relative to the setting with material payoffs only. However, these effects are fragile and do not withstand competition between sellers, which is welfare-decreasing.

We explore in-depth strategies used by sellers and buyers to understand how these aggregate outcomes emerge. The sellers with low-quality goods lie less in the markets with psychological payoffs and without competition relative to the other two markets. In particular, when psychological payoffs are present, sellers lie more in the markets with competition. This translates into how informative messages are in the two markets and how they influence buyers' purchasing decisions. Interestingly, even though buyers correctly interpret messages in markets without competition, they fail to account for the increased dishonesty of sellers in markets with competition. The combination of higher dishonesty among the low-quality sellers and buyers' inability to factor in this dishonesty results in a decrease in welfare for all market participants when there is competition in markets with psychological payoffs.

Later in the paper, we explore mechanisms behind competitive forces promoting immoral behavior (dishonesty) by sellers in markets with psychological payoffs. We document that sellers respond to each other by lying more often when they observe another seller lying more often. This dynamic is consistent with the fear of being excluded from interaction with the buyer, who tends to choose sellers who claim to have a high-quality product. Furthermore, we show that such a response is optimal when sellers are unsure of the strategy their opponent is playing. In other words, the strategic uncertainty about other sellers' actions, which is present both in actual and our laboratory markets, pushes the low-quality sellers towards more lying in order to win the competition for a single buyer.

As for the buyers, we use additional belief data collected in the experiment to investigate why buyers struggle to grasp the correct interpretation of messages, often perceiving them as more informative than they truly are, in markets with competition. In contrast, such mistakes occur far less often in markets without competition. We consider various explanations, some of which are rooted in cultural factors, drawing from the widespread belief that competition between sellers typically benefits consumers, and others are driven by the feedback available to buyers in the two markets. To support the latter perspective, we present empirical evidence of the sluggishness of beliefs in the markets with competition and discuss the theoretical forces that lead to this conservative belief updating in such markets compared to those without.

Contribution and Connection to the Literature. Our theoretical model builds on two branches of literature: psychological games that incorporate belief-dependent preferences (Geanakoplos et al., 1989; Battigalli and Dufwenberg, 2007, 2009, 2022) and games with lying costs (Chen et al., 2008; Kartik, 2009; Sobel, 2020). We focus on three psychological forces, which have been identified in the literature as the leading factors in communication games with hidden actions and hidden information: lie aversion, guilt aversion, and disappointment aversion (Gneezy, 2005; Hurkens and Kartik, 2009; Charness and Dufwenberg, 2006, 2011; Vanberg, 2008; Goeree and Zhang, 2014; Casella et al., 2018; Abeler et al., 2019). Relative to the above-mentioned papers, the contribution of our model is to explicitly model these forces, derive equilibrium predictions for the game they define with and without competition, and document the multiplicity of equilibria that emerge in such a setting.

Our experiment is inspired by the seminal papers of Charness and Dufwenberg (2006, 2011) but uses a design that is conceptually very different. Charness and Dufwenberg (2011) study a different version of a hidden information game and test for the *presence* of psychological forces in communication games by inducing only material payoffs and observing outcomes different from those predicted by the material-payoff-only model. By contrast, as discussed above, we use the induced value approach, which has been successfully implemented in a variety of individual decision-making tasks and strategic settings but has not yet been used in games with psychological payoffs. We see our implementation of this approach as one of the contributions of this paper.

The experimental literature concerning the interplay between competition and communication is still in its infancy.[6] The three most closely related papers to ours are Casella et al. (2018), Goeree and Zhang (2014), and Born (2020). Casella et al. (2018) study a communication game with hidden actions and communication among competing senders but do not model the game as a psychological game. The authors find that messages are inflated in the game with competition, but these inflated messages induce mostly the same actions from receivers, indicating that the receivers account for this inflation. Our game is the game with hidden information rather than hidden actions, and our results reveal different patterns: as in Casella et al. (2018), we find a shift in the communication strategies when competition is present, but contrary to Casella et al. (2018), our buyers fail to interpret messages correctly when competition is present. Instead, buyers in our experiment behave as if they believe messages have the same meaning in the presence as well as in the absence of competition.[7]

Closer to our setup, Goeree and Zhang (2014) introduce competition in the hidden-information game studied in Charness and Dufwenberg (2011). They find competition and communication act as substitutes. Communication raises efficiency in the absence of competition but lowers efficiency when competition is present. Similarly, competition raises efficiency without communication but lowers it when parties can communicate with each other. The authors briefly discuss some behavioral explanations that can account for such outcomes, including inequality aversion, guilt aversion, lying aversion, and reciprocity. Although our paper shares some of the features of Goeree and Zhang (2014) with respect to the way we define material payoffs and competition, we take a very different approach by modeling the game as a psychological game in which players exhibit a wide range of emotions (translated into their payoffs). We then obtain theoretical results regarding the effects of competition on market outcomes and players' behavior and test these predictions in a lab experiment in which we induce payoffs associated with these emotions. Despite different approaches, both Goeree and Zhang (2014) and we show that competition decreases efficiency in a game with communication.

Born (2020) studies promise competition between sellers who differ in their intrinsic motivation and costs of breaking promises. This model features both hidden information and hidden actions of sellers. Theoretically, Born shows that, on average, promise competition increases buyers' welfare relative to a no-competition case, because some sellers promise more than they would in the absence

[6]Several studies look at the effects of competition on trust in various environments (Huck et al., 2012; Keck and Karelaia, 2012; Fischbacher et al., 2009). See also Vespa and Wilson (2016), who study experimentally a multi-dimensional communication game with multiple senders and find that in this very different setting, receivers do not use the information optimally.

[7]Similar results are found by Jin et al. (2021) when studying disclosure behavior by sellers in a market. In that market, a failure to disclose the quality of one's product should signal its low quality, yet buyers fail to completely adjust to it.

of competition. Experimental results reveal that sellers' behavior crucially depends on their game experience as the difference between the competition and the no-competition case was observed only in the first rounds of the experiment. Contrary to Born's results, we observe significant welfare differences between the game with and without competition after subjects have learned to play the game and have converged to stable behavior.

Structure of the Paper. We proceed as follows. In Section 2, we introduce the three games and solve for their equilibria. In Section 3, we describe the experimental design and its implementation. Section 4 contains the results of the experiment, while Section 5 investigates the behavioral mechanisms underlying our results. Section 6 offers some conclusions.

2 The Model

In this section, we present three variants of the communication model that we study in the lab: (1) the game with only material payoffs, (2) the game with both material and psychological payoffs and without competition, and (3) the game with both types of payoffs and competition between sellers. The detailed analysis of the last two games is presented in Section 1 in the Online Appendix.

2.1 The Game with Material Payoffs and without Competition

Setup. We study a communication game between an informed seller (he) and an uninformed buyer (she). The seller owns one unit of the product and wants to sell it to the buyer. The product can be either low quality, $q = q_L$, with probability $p > \frac{1}{2}$, or high quality, $q = q_H$, with remaining probability $1-p$. The seller knows the quality of the product he owns and sends a message, m, to the buyer in an attempt to convince her to purchase his good. Two messages are possible: m_1="The product is really high quality" and m_0 ="The product is low quality." $M = \{m_0, m_1\}$ denotes the set of possible messages. The buyer does not know the quality of the good but observes the message sent by the seller. The buyer is interested in purchasing the high-quality product and not the low-quality product. The situation is complicated by the fact that the buyer cannot distinguish the high- from the low-quality good until she purchases it and has to instead rely on the seller's messages. After observing the message, the buyer either buys, or does not, and the game ends.

The material payoffs of players are depicted in Figure 1. When a good is not sold, the buyer and the seller each receive a fixed payoff of 5. When a high-quality good is sold, both receive a payoff of 10. The interesting case arises when the seller manages to peddle off a low-quality good: in this case, the seller receives a payoff of 21, while the buyer receives 0. Because in this case, the preferences of the buyer and the seller are misaligned, the potential for lying exists.[8]

Equilibria. Any Bayesian Nash equilibrium outcome in this game features no trade. To see why, assume by contradiction that an equilibrium exists in which, after observing message m_i, the buyer purchases the product with a higher probability than after observing a message m_j. Such behavior is justified if the buyer believes the seller with a high-quality product is more likely to send message m_i than message m_j. However, in that case, the seller with a low-quality product will mimic this behavior and will also send a message m_i, which contradicts our initial presumption. Thus, no

[8]We use the term lying to refer to the situation in which the seller's message does not match the product quality he owns, i.e., when a low-quality seller sends the m_1 message and when a high-quality seller sends the m_0 message.

Figure 1: Payoffs in the Game with Material Payoffs and without Competition

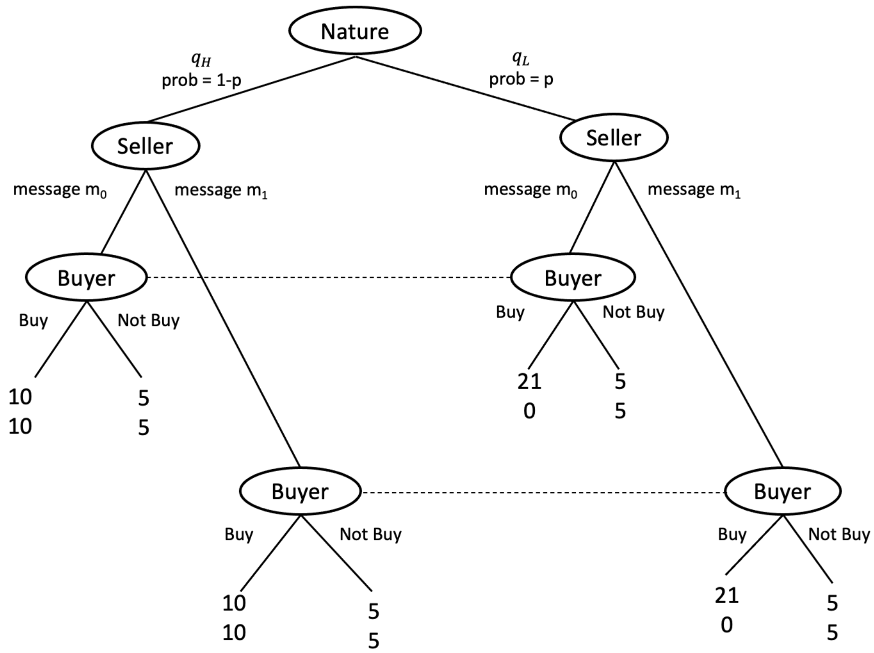

Notes: At each node, the top payoff depicts the seller's payoff, while the bottom one depicts the buyer's payoff. The dashed line indicates the buyer's information set because she does not know the type of seller she is dealing with.

equilibrium can exist in which one message entails a higher probability of a high-quality product than another. Therefore, the buyer is left with her prior beliefs, and given the material payoffs, *no trade is the only equilibrium outcome.*[9]

2.2 The Game with Psychological Payoffs and without Competition

The situation changes when we introduce psychological payoffs. Players are now motivated not only by their material payoffs but also by belief-dependent utilities which are determined by players' strategies and their beliefs.

Setup. There are two players: the seller (he) and the buyer (she). The seller owns one unit of the product and wants to sell it to the buyer. The seller sends a message m to the buyer in an attempt to convince her to purchase the good.

The seller's type $t^S \in T^S$ consists of three elements: the product quality $q \in \{q_L, q_H\}$, the guilt sensitivity $g \in \{0, G\}$, and the lying sensitivity $l \in \{0, L\}$, where $q_L < q_H$, $G > 0$, and $L > 0$. Thus,

[9]There are at least two ways to sustain a no-trade equilibrium outcome in our game: one, in which all sellers send message m_0 and another in which all sellers send message m_1.

there are eight possible types of the seller:

$$T^S = \{(q_L, 0, 0), (q_L, G, 0), (q_L, 0, L), (q_L, G, L), (q_H, 0, 0), (q_H, G, 0), (q_H, 0, L), (q_H, G, L)\}.$$

The buyer's type $t^B \in T^B$ is captured by a single disappointment sensitivity parameter ω.

At the start of the game, nature draws the types of the seller and the buyer independently. The buyer type is drawn according to a uniform distribution with support equal to the unit interval $[0, 1]$, i.e., $\omega \sim U[0, 1]$. The seller's type is drawn from the following distribution. The first four seller types, in which the quality is low, occur with probability $\frac{p}{4}$, while the last four types, in which the quality is high, occur with probability $\frac{1-p}{4}$. Thus, as in the game without psychological payoffs, the product has a low quality with probability p and a high quality with the remaining probability $1 - p$. For convenience, we also define $T^S_{q_H} \subseteq T^S$ as the subset of seller types with high-quality products, and $T^S_{q_L} \subseteq T^S$ as the subset of seller types with low-quality products.

After nature draws the buyer and seller types, the seller observes his type t^S, but not the buyer's type, while the buyer observes her type t^B, but not the seller's type. The distributions from which nature draws the buyer and seller types are common knowledge.

Next, the seller chooses one of two messages to send to the buyer before the buyer chooses whether to buy or not. The seller's strategy maps his type into (possibly randomized) messages he sends to the buyer, that is, $s^S : T^S \to \Delta(M)$ where $M = \{m_0, m_1\}$ and $\Delta(M)$ represents the space of probability distributions over messages M.

The buyer's strategy maps her type and the message she receives into a (possibly randomized) purchasing decision, that is, $s^B : T^B \times M \to \Delta(\{\text{Buy}, \text{Not Buy}\})$.

As in the game with material payoffs only, the messages have a natural meaning, i.e., m_0 corresponds to "The product is low quality" and m_1 corresponds to "The product is really high quality". This meaning is important for defining the psychological parts of players' payoffs as we describe below. We let $p^B(m_i|t^S)$ denote the probability that the buyer thinks a seller of type t^S sends the message m_i. Given these beliefs, the buyer updates on the probability that the product is high quality using Bayes' rule:

$$z^B(m_i) = \Pr[q = q_H|m_i] = \frac{(1-p) \cdot \sum_{t^S \in T^S_{q_H}} \frac{1}{4} p^B(m_i|t^S)}{(1-p) \cdot \sum_{t^S \in T^S_{q_H}} \frac{1}{4} p^B(m_i|t^S) + p \cdot \sum_{t^S \in T^S_{q_L}} \frac{1}{4} p^B(m_i|t^S)}. \tag{1}$$

We refer to $z^B(m_i)$ as the buyer's interpretation of message m_i and note that it simply represents the probability that, anticipating the messages different sellers will send, the message m_i comes from a high-quality seller.

As we describe below, the disutility from guilt requires the seller to formulate beliefs over the buyer's beliefs so that conditional on each message the seller can calculate the expected payoff of the buyer. Thus, for any $m_i \in M$, we let $p^S(m_i|t^S)$ denote the probability that the seller places on the buyer believing he sends a message m_i conditional on being of type t^S. We let $z^S(m_i)$ denote the probability that, according to the seller's beliefs, the buyer thinks that message m_i comes from a high-quality seller and calculate $z^S(m_i)$ using Bayes' rule as

$$z^S(m_i) = \Pr[(q_H, \cdot, \cdot) \text{ sends } m_i] = \frac{(1-p) \cdot \sum_{t^S \in T^S_{q_H}} \frac{1}{4} p^S(m_i|t^S)}{(1-p) \cdot \sum_{t^S \in T^S_{q_H}} \frac{1}{4} p^S(m_i|t^S) + p \cdot \sum_{t^S \in T^S_{q_L}} \frac{1}{4} p^S(m_i|t^S)}. \tag{2}$$

We refer to $z^S(m_i)$ as the seller's interpretation of the buyer's interpretation of message m_i.

Payoffs. The players' payoffs consist of two parts: the material payoffs, which depend on the quality of the seller's product and the buyer's purchasing decision, and the psychological payoffs, which depend on the psychological disutilities players may experience in our game. The material part of players' payoffs are denoted by Π^B and Π^S for the buyer and the seller, respectively, and are the same as in the game without psychological payoffs, i.e.,

$$\Pi^B(q_H, \text{Not Buy}) = \Pi^S(q_H, \text{Not Buy}) = \Pi^B(q_L, \text{Not Buy}) = \Pi^S(q_L, \text{Not Buy}) = 5,$$

$$\Pi^B(q_H, \text{Buy}) = \Pi^S(q_H, \text{Buy}) = 10, \quad \Pi^B(q_L, \text{Buy}) = 0, \quad \text{and } \Pi^S(q_L, \text{Buy}) = 21.$$

The psychological part of the buyer's payoff is determined by her **disappointment** sensitivity. The buyer may experience disappointment whenever she relies on the seller's false claims of a high-quality product and ends up purchasing the product. Thus, the disappointment disutility kicks in only if she purchases the low-quality product and is proportional to the difference between her expected and actual material payoff, and is weighted by her disappointment sensitivity ω, i.e.,

$$\text{Disappointment}(\omega) = \omega \cdot \left(10z^B(m_i) - 0\right).$$

The buyer's total expected payoff from purchasing the good after receiving the message m_i is

$$\mathbb{W}^B\left(\cdot|m_i, \text{Buy}\right) = 10z^B(m_i) - \omega \cdot (10z^B(m_i) - 0) \cdot (1 - z^B(m_i)). \tag{3}$$

The psychological part of the seller's payoff consists of guilt aversion and lying aversion.

The seller's disutility from **guilt** comes from disappointing buyers who are sensitive to disappointment. A seller may feel guilty for leading on the buyer (by claiming he has a high-quality product even though he does not) and eventually delivering the low-quality product. All else being equal, the higher the disappointment parameter ω of the buyer, the more guilty the seller feels when he leads the buyer on and then sells her the low-quality product. Formally, we model the seller's guilt as a cost that is proportional to the difference between the buyer's expected payoff from purchasing the good, and her actual payoff, i.e., $(10 \cdot z^S(m_i) - 0)$, where $10 \cdot z^S(m_i)$ represents the seller's belief regarding the expected payoff of the buyer from buying the product after observing m_i, and 0 represents the actual material payoff the buyer receives. The guilt is also proportional to the seller's expectation of ω conditional on the buyer buying. So, the amount of guilt that the seller with the guilt-sensitivity parameter g experiences when he sends message m_i is

$$\text{Guilt}(m_i, g) = g \cdot \left(10z^S(m_i) - 0\right) \cdot \mathbb{1}(m_i, \text{Buy}) \cdot \mathbb{E}[\omega|m_i, \text{Buy}],$$

where $\mathbb{1}(m_i, \text{Buy})$ is an indicator function taking value 1 if the buyer purchases the product after receiving the message m_i and value 0 otherwise.

Our definition of guilt is reminiscent of the notion of deception in Sobel (2020), where deception depends on how the receiver interprets messages and how her actions might change in response to them.[10] It also relates to Battigalli and Dufwenberg (2007) concept of guilt in games, which captures a failure to live up to receivers' anticipated behavior.

[10] In the Sobel (2020) framework, the deception and lying capture two very different notions because one player can deceive another without lying, and not all lies constitute a deception.

In addition, the seller may experience disutility from **lying** if he sends an untruthful message weighted by his lying sensitivity l. This definition is consistent with theoretical notions used in Kartik (2009) and Sobel (2020) as well as experimental evidence surveyed by Abeler et al. (2019).[11] As Sobel (2020) notes, the definition of a lie depends on the existence of accepted meanings of words. This is exactly what we do in our paper: sellers' messages have precise meanings rather than context-free neutral labels. In the analysis that follows (and in the experiment), we abstract away from penalizing the sellers with high-quality products who lie and send the message m_0. Although sending an m_0 message when $q = q_H$ is indeed a lie, it is a self-destructive one.

Collecting all the terms, we obtain the seller's total expected payoff

$$\mathbb{W}^S(\cdot) = \begin{cases} \Pi^S(q_H, s^B) & \text{if } q = q_H \\ \Pi^S(q_L, s^B) - g \cdot 10z^S(m_i) \cdot \mathbb{E}[\omega|m_i, \text{Buy}] \cdot \mathbb{1}(m_i, \text{Buy}) - l \cdot \mathbb{1}(m_i = m_1) & \text{if } q = q_L \end{cases}, \quad (4)$$

where $\mathbb{1}(m_i = m_1)$ is an indicator function taking value 1 if the seller sends the message m_1 and value 0 otherwise.

Timing of the Game. The nature draws the types of the seller and the buyer. Given his type, the seller sends a message to the buyer. Given her type and the message she receives, the buyer chooses whether to purchase the product or not. The payoffs are realized.

Informative Equilibria. The introduction of psychological payoffs introduces the possibility of sustaining equilibria in which some information about the product quality is conveyed in the communication stage. As we discuss in the Online Appendix, we make some restrictions on the parameter space, which are satisfied in our experiment, and look for perfect Bayesian equilibria of this game that are in pure strategies, in which messages have their intended meaning such that sellers with the high-quality product send the message m_1.

To illustrate the trade-offs faced by the players, consider first the decision of the buyer. The buyer prefers to purchase the product when

$$10z^B(m_i) - \omega \cdot (10z^B(m_i) - 0) \cdot (1 - z^B(m_i)) \geq 5.$$

This inequality defines the threshold value $\bar{z}^B(m_i, \omega) \in (0, 1)$ such that the buyer with disappointment sensitivity ω wants to buy the product if and only if $z^B(m_i) \geq \bar{z}^B(m_i, \omega)$. The threshold $\bar{z}^B(m_i, \omega)$ is increasing in ω, so buyers who are more sensitive to disappointment are less inclined to buy. Moreover, for each value of $z^B(m_i)$, there exists a threshold value $\bar{\omega}(m_i)$, such that only buyers with $\omega < \bar{\omega}(m_i)$ purchase the product. This threshold is $\bar{\omega}(m_i) = \min\{\frac{2z^B(m_i)-1}{2z^B(m_i)(1-z^B(m_i))}, 1\}$ if $z^B(m_i) \geq \frac{1}{2}$ and $\bar{\omega}(m_i) = 0$ otherwise, respecting the fact that $\omega \in [0, 1]$.

Consider now the seller's behavior. All high-quality sellers send the message m_1, thus if the buyer observes message m_0, she knows it comes from the low-quality seller, that is, $z^B(m_0) = 0$. Thus, if in equilibrium there is some trade, it must occur when the buyer receives the message m_1. So, in any equilibrium that supports trade, the seller with the low-quality product and no guilt and no lying aversion strictly prefers to send the message m_1. What about other low-quality sellers who

[11] Abeler et al. (2019) provide a meta-data analysis of experimental work on lying and show that people have a preference to be seen as honest and a preference for honesty per se which explains why people lie less than what theory predicts they should if they only cared about material payoffs.

experience at least some psychological costs? The type (q_L, g, l) will send a truthful message m_0 if and only if

$$5 \geq \bar{\omega}(m_i) \cdot \left(21 - g \cdot 10 z^S(m_1) \cdot \frac{\bar{\omega}(m_1)}{2} - l\right) + (1 - \bar{\omega}(m_1)) \cdot (5 - l)\,, \tag{5}$$

where $\bar{\omega}(m_i) = \Pr[\omega \leq \bar{\omega}(m_i)]$ is the probability that a buyer will purchase the product after receiving message m_i and $\frac{\bar{\omega}(m_i)}{2}$ is the expected disappointment of the buyer who purchases the product after receiving message m_i. The restrictions we impose on the parameter space guarantee that the low-quality seller who suffers from lying aversion, i.e., $l = L > 0$, will refrain from lying even if he does not suffer from guilt aversion, i.e., $g = 0$. The optimal behavior of the remaining seller type $(q_L, G, 0)$ depends on which equilibrium is played. The multiplicity of equilibria arises because guilt-averse sellers of low-quality products suffer a larger disutility from selling to more optimistic buyers.[12]

Given the discussion above, it follows that in any informative equilibrium $0 = z^B(m_0) < 1 - p \leq z^B(m_1) < 1$. We call perfect Bayesian equilibria with these properties **Partially Informative Equilibria (PIE)**.

2.3 The Game with Psychological Payoffs and with Competition

In the game with competition, the buyer faces two sellers. After learning their private types, each seller sends a message to the buyer about the quality of his good. Upon observing the two messages, the buyer selects one of the sellers to deal with and then chooses whether to purchase the product from the selected seller or not.

Setup. This game has three players: the buyer, seller S_1, and seller S_2. Each seller has one of the eight possible types determined by the combination of his product quality $q^{S_j} \in \{q_L, q_H\}$, his guilty sensitivity $g^{S_j} \in \{0, G\}$, and his lying sensitivity $l^{S_j} \in \{0, L\}$. Thus, the seller S_j's type is $t^{S_j} \in T^S$. The buyer's type $t^B \in T^B$ is captured by a disappointment parameter ω. At the start of the game, nature draws the types of the buyer and two sellers independently from the same distributions as in the game without competition. The buyer knows her own type but not the types of the sellers. Each seller knows his own type but not the type of the other seller or the buyer.

The seller S_j's strategy maps his type into (possibly randomized) messages, i.e., $s^{S_j} : T^S \to \Delta(M)$, where $M = \{m_0, m_1\}$ and $\Delta(M)$ represents the space of probability distributions over messages M. The messages have the same pre-specified meaning as in the game without competition.

The buyer's strategy consists of two decisions: the selection decision and the purchasing decision. The selection decision maps the two observed messages into the (possibly randomized) choice over the two sellers, that is, $k^B : M \times M \to \Delta(\{S_1, S_2\})$ where $\Delta(\{S_1, S_2\})$ represents the space of probability distributions over two sellers. In some of the analyses, it will be useful to refer to the selected seller as S^{win} and the one that lost as S^{lose}. The purchasing decision maps the buyer's type and the message of the selected seller into a (possibly randomized) purchasing decision, that is, $s^B : T^B \times M \to \Delta(\{\text{Buy}, \text{Not Buy}\})$.[13]

[12] So when buyers don't expect this type to sell to them it is the best response for this type to not sell to them, but when these buyers do expect this type to sell to them it is the best response for this type to sell to them.

[13] We separate the buyer's decision into two steps for clarity and comparability between the game with and without competition. Note that there is no new information that the buyer receives between the selection decision and the purchasing decision, which is why she bases her purchasing decision on the original message of the selected seller.

We let $p^B(m_i|t^{S_j})$ denote the probability that the buyer thinks the seller S_j with type t^{S_j} sends the message m_i. Given these beliefs, the buyer updates on the probability that the product owned by the seller S_j is of high quality using Bayes' rule. The calculation is similar to the one described in equation (1) except that the buyer updates on each seller separately.[14]

Since some seller types experience disutility from guilt, each seller is required to formulate beliefs over the buyer's beliefs in case he wins the competition, so that conditional on each message the expected payoff of the buyer can be calculated. For any $m_i \in M$, we let $p^{S_j}(m_i|t^{S_j})$ denote the probability that the seller S_j thinks the buyer thinks he sends message m_i when his type is t^{S_j}. We then use Bayes' rule to calculate $z^{S_j}(m_i)$, which is the probability that according to the seller S_j's beliefs, the buyer thinks that message m_i comes from a high-quality seller S_j (see equation (2)).

The buyer's total expected payoff from purchasing the good after receiving the message $m^{S^{\text{win}}} = m_i$ is depicted by equation (3). The non-selected seller gets zero, i.e., $\mathbb{W}^{S^{\text{lose}}} = 0$. The expected payoff of the selected seller who sent $m^{S^{\text{win}}}$ is defined as in (4), i.e., $\mathbb{W}^{S^{\text{win}}} = \mathbb{W}^S(m_i = m^{S^{\text{win}}})$.

Timing of the Game. The nature draws the types of all three players, the two sellers and the buyer. Types are privately observed by players. Each seller sends a message to the buyer. The buyer observes both messages simultaneously and chooses one seller to deal with in the remainder of the game. Then, the buyer chooses whether to purchase the product from the selected seller or not. The payoffs are realized and the game is over.

Informative Equilibria. We look for Perfect Bayesian equilibria that support trade in the above game with parameters implemented in the experiment, in which the buyer's purchasing decisions and the sellers' communication strategies are pure, messages have their intended meaning such that sellers with the high-quality product send the message m_1, and in which the buyer resolves indifference about which seller to select by flipping a coin. We focus on symmetric equilibria in which both sellers use the same communication strategy and, as before, refer to such equilibria as **Partially Informative Equilibria (PIE)**.

2.4 Parameterization

To bring our model to the lab, we need to set parameters for the games. Our selection of parameters is guided by two principles. First, we opted for the probability of a low-quality product to be $p = 60\%$, ensuring that, in the absence of psychological payoffs, the unique pooling equilibrium outcome is one in which the buyer never purchases the product. Second, we set $G = 6$, $L = 20$, and $\omega \sim U[0, 1]$ to ensure that the set of equilibria supported in the two games with psychological payoffs is identical. This set consists of three equilibria:

1. **Pooling equilibrium.** All types of sellers send message m_0. The buyer treats messages as uninformative and does not update her prior beliefs about the product quality, regardless of the observed message. That is, after observing either message, the buyer believes that the chance she is facing a low-quality seller is 60%. In the game with competition, the buyer randomly selects one seller. The buyer does not purchase the product and collects a payoff of

[14] In the game without competition, the buyer's beliefs induce a pair of probabilities, one for each message, i.e., $(z^B(m_0), z^B(m_1))$. In the game with competition, the buyer forms such a pair for each seller, i.e., $(z^B_{S_1}(m_0), z^B_{S_1}(m_1))$ and $(z^B_{S_2}(m_0), z^B_{S_2}(m_1))$, where $z^B_{S_j}(m_i)$ denotes the probability that the buyer believes the message m_i sent by seller S_j comes from the high-quality seller S_j.

5. The seller gets a payoff of 5 in the game without competition and an average payoff of 2.5 in the game with competition.[15]

2. **PIE1.** The seller types $(q_H, \cdot, \cdot)$, $(q_L, 0, 0)$ and $(q_L, G, 0)$ send message m_1, while the remaining types send message m_0. In our parameterized game with competition, the buyer selects a seller with message m_1 if she receives two different messages; otherwise, she randomly selects one seller. If the message of the (selected) seller is m_1, the buyer believes that there is a 57% chance that this message comes from the high-quality seller and only buyers with $\omega \leq 0.29$ buy the product. Thus, the buyer purchases the good with a probability of 0.29 after receiving an m_1 message from the (selected) seller. If, however, the (selected) seller's message is m_0, the buyer knows this message is sent by the low-quality seller and does not buy the product. The buyer's expected payoff is 5.07 in the game without competition, and 5.09 in the game with competition. The sellers' expected payoffs are 8.05 in the game without competition, and 3.65 in the game with competition.

3. **PIE2.** The seller types $(q_H, \cdot, \cdot)$ and $(q_L, 0, 0)$ send message m_1, while all the remaining types send message m_0. In the game with competition, the buyer selects a seller with message m_1 if she receives two different messages; otherwise, she randomly selects one seller. If the message of the (selected) seller is m_1, the buyer believes that there is a 73% chance that this message comes from the high-quality seller and all buyer types purchase the product. Thus, the buyer purchases the good with a probability of 1 after receiving an m_1 message from the (selected) seller. If, however, the (selected) seller's message is m_0, the buyer knows this message is sent by the low-quality seller and does not buy the product. The buyer's expected payoff is 5.70 in the game without competition, and 6.02 in the game with competition. The sellers' expected payoffs are 13.80 in the game without competition, and 5.69 in the game with competition.

The three equilibria described above are ranked in terms of how much information sellers transmit in the communication stage. We define the informativeness of an equilibrium as the difference between buyers' posterior beliefs after observing the two messages, i.e., $Eq^{\text{info}} = z^B(m_1) - z^B(m_0)$. The larger this difference, the more information the buyer learns from the sellers' messages. Then, the least informative equilibrium is the pooling one, while the most informative one is the PIE2, in which only the low-quality sellers with no psychological costs lie in the equilibrium, i.e.,

$$0 = Eq^{\text{info}}_{\text{POOL}} < Eq^{\text{info}}_{\text{PIE1}} = 0.57 < Eq^{\text{info}}_{\text{PIE2}} = 0.73.$$

Informativeness of equilibria directly translates into buyers' expected payoffs: the more information the buyer receives from the seller's messages, the better purchasing decisions she can make. This can be seen by comparing the buyers' expected payoffs across three equilibria holding fixed the presence or absence of competition between sellers. Finally, we note that if the same PIE is played in both games with and without competition, then buyers benefit from competition and earn higher expected payoffs. This happens because the presence of two sellers increases the likelihood that the buyer will select a high-quality seller to deal with. However, this observation relies on the assumption that the same equilibrium is played in the two games. Whether such an assumption is reasonable or not is ultimately an empirical question, which we address next in our experiment.[16]

[15] This pooling equilibrium is quite fragile as it requires all sellers to send message m_0 to utilize the fact that sending message m_0 is a free lie for a high-quality type.

[16] Schotter et al. (1996) have shown competition can have an impact on rejection behavior in ultimatum games in which receivers are more willing to accept low offers if the person making such an offer had to compete in a tournament-like setting.

2.5 Discussion of Modeling Choices

Before we turn to the experiment, let us discuss some of our modeling choices and their relation to the existing literature.

Game tree and material payoffs. The structure of our game is closely related to the communication game of Crawford and Sobel (1982) in which an informed sender (the seller) sends a message to an uninformed receiver (the buyer), who then takes an action that affects both players' payoffs. We chose to focus on a simpler version of such a game with senders (sellers) possessing high- or low-quality products and receivers (buyers) having to decide whether to buy.

The important feature of our setting is that a buyer knows neither the quality of the seller's product nor the psychological type of the seller she is dealing with; that is, our game belongs to the class of communication games with hidden information. One of the key experimental papers in this literature is that of Charness and Dufwenberg (2011), CD-11 hereafter. Although our game shares some similarities with one of the games studied in CD-11, important differences also exist.[17] First, in the CD-11 game, the seller can choose not to trade with the buyer even if the buyer wants to trade; in this case, both get a fixed no-trade payoff. This adds an element of reciprocity on the part of the informed player (the seller). By contrast, our game is a pure communication game in which a seller can only send a message to a buyer, and otherwise has no action to take. Second, in the CD-11 game, only high-quality-product sellers prefer to trade, whereas the low-quality-product sellers prefer a fixed no-trade payoff. Our game differs, in that all sellers want to trade irrespective of the quality of their product. Third, the CD-11 game has an additional element that is absent in our game, namely, a positive probability that trade may be prevented from occurring even if both parties agree to trade. This feature adds another layer of uncertainty and reduces the buyer's ability to infer the seller's product quality even at the end of the game. We chose to study what we feel is the simplest communication game with hidden information, which is amenable to psychological payoffs and the introduction of competition among sellers.

The role of psychological payoffs. One may wonder whether a psychological game is at all necessary to study the effect of competition in our communication game. In other words, what would happen in a standard cheap-talk game, in which sellers may experience aversion to lying but no guilt aversion, and buyers do not feel disappointed when they expect outcomes that do not materialize? Such a game would be the standard game since sellers' guilt and buyers' disappointment are what make our game a psychological game, i.e., the game in which payoffs depend on players' beliefs. Depending on the distribution of sellers' types, such a game will feature at most one informative equilibrium, in which all high-quality sellers are truthful and some low-quality sellers send the untruthful message m_1.[18]

[17] Here, we focus on the one game studied in CD-11, which is closest to our game in the sense that the low-quality seller can gain materially from trading with the buyer if and only if he fools the buyer into buying his product, which the buyer would prefer not to buy. The authors also study another version of the communication game in which the low-quality seller can benefit materially from trading with the buyer, even if the buyer knows she is buying a low-quality product. The introduction of communication between the seller and the buyer is more effective in the second game than in the first one.

[18] Kartik (2009) analyzes the general version of such a model without competition between sellers and finds that partially informative equilibria exist, in which sellers with higher types pool together, while lower types separate. These equilibria are characterized by inflated language in the sense that sellers might claim they have a higher type than they actually have.

The presence of both guilt aversion and lying aversion is what allows us to sustain *two different informative equilibria* in both games with and without competition. The multiplicity of informative equilibria opens up the possibility of competition being welfare decreasing via the selection of the welfare-inferior but still somewhat informative equilibrium. To see this point, re-examine the inequality (5) that guarantees that a seller with type (q_L, g, l) sends truthful message m_0 in the game with psychological payoffs but without competition. Re-arranging the terms, we get

$$l + g \cdot 10z^S(m_1) \cdot \frac{\bar{\omega}(m_1)^2}{2} \geq 16 \cdot \bar{\omega}(m_1).$$

Both guilt and lying aversion prevent some low-quality sellers from lying to a buyer. However, the two forces play out differently. The first term on the left captures that low-quality sellers with a high aversion to lying per se are likely to be truthful in the communication stage because their own intrinsic penalty for lying is too high. The second term on the left captures the disutility from lying resulting from the guilt that one experiences about misleading the buyer; the extent of sellers' guilt depends on the sellers' beliefs about buyers' interpretation of message m_1. Either of the two channels can prevent low-quality sellers from sending an untruthful message m_1. But, the presence of *both* channels is what allows us to sustain two different informative equilibria, which can be ranked in terms of how much information is contained in message m_1. Absent these psychological forces, our game would be considerably less interesting.

3 Experimental Design

The experiment was conducted in the experimental lab of the Center for Experimental Social Science (CESS) at New York University. We recruited 179 subjects via E-mail from the general undergraduate population at NYU for an experiment that lasted approximately one hour and 45 minutes. Subjects received a show-up fee of $7 and on average received a final payment of $29.50 for their participation. The program used in the experiment was written in Z-Tree (Fischbacher (2007)). We present our experimental design and treatments' variation in Section 3.1. In Section 3.2 we discuss the benefits and challenges of inducing psychological costs in the lab experiment and describe how we deal with eliciting both subjects' actions and subjects' beliefs.

3.1 The Design

Our experiment is a direct implementation of the three games described above. We conducted three separate treatments: a Monetary treatment, a No Competition treatment, and a Competition treatment. Each experimental session consisted of only one of the three treatments. Once in the lab, subjects were randomly assigned to play the role of either a buyer or a seller, and these roles remained fixed during the entire session. We refer the reader to Section 2 in the Online Appendix for the complete set of instructions in one of the treatments and describe below the main features of the experimental protocol.

In the **No Competition treatment**, the subjects play the communication game described in Section 2.2 with the parameters described in Section 2.4. Specifically, the seller's task is to specify a decision function that maps his type (q, g, l) into a message from the set of messages $\{m_0, m_1\}$. The sellers enter their decisions by filling out a table with eight entries: one for each type they might be assigned.

The buyer's task is to enter a purchasing decision conditional on the message she receives and her sensitivity type ω. Buyers do that in the experiment by entering two cutoff values, $\omega'(m_0)$ for message m_0, and $\omega'(m_1)$ for message m_1, such that whenever the realized value of ω is less than the cutoff value for the received message, the buyer buys the good. For values of ω above the cutoff value, the buyer does not buy the good.

In addition to specifying their strategies, subjects are also asked to enter their beliefs. Each buyer is asked to enter a number between 0 and 100 representing her belief that the sellers who sent the message m_i possessed a high-quality good. They did so for both messages m_0 and m_1. Each seller was asked to enter a number representing his (second-order) belief about what the seller thought the first-order belief of the buyer was, upon receiving either message m_0 or m_1.

Once the subjects had specified their strategies and beliefs, these choices were simulated for 10 periods, where the computer randomly determined the quality of the good and a type for each seller, and a sensitivity parameter for each buyer for every period. Further, using the strategies they entered, the computer determined payoffs for them for each of the 10 periods. We call these 10 periods *a block*, and each treatment had 10 such blocks. After each block, the subjects were given time to review their actions and payoffs for the preceding 10 periods before entering their strategies and beliefs again for the next block that determined their payoffs for the next 10 periods. In each block, subjects maintained their roles but were randomly assigned a new partner.[19]

We use this block design because entering a strategy, a set of beliefs, and reviewing feedback is a time-consuming process, and, hence, it would be practically impossible for subjects to do that for, say, 50 periods. Our design allows subjects to maximize the amount of feedback they get while economizing on the time they spend mechanically entering their strategies and beliefs. More importantly, we feel this approach is beneficial when one conducts experiments using the strategy method, because once a strategy is entered, one might as well receive a lot of feedback on it before being asked to change it.[20] Entering a strategy and receiving only one period of feedback does not allow a subject to learn very much about it.

To determine a subject's payoff in the experiment, we randomly chose one of the 10 blocks, and in that block paid subjects either for their payoffs in the game or for their elicited beliefs, using a quadratic scoring rule (for a similar approach, see Nyarko and Schotter (2002)). Eliciting both actions and true beliefs in psychological games is tricky because payoffs are a function of beliefs. We discuss this issue in detail in Section 3.2 and describe how we dealt with it.

Finally, at the end of the session, we administered two risk-elicitation tasks using the Gneezy and Potters (1997) methodology. In each of these two tasks, we asked subjects to allocate 200 points (translating into \$2) between a safe investment, which had a unit return (i.e., returning point for point), and a risky investment, which with probability p returned R points for each point invested and with probability $1 - p$ produced no returns for the investment. In the first task, $p = 0.5$ and $R = 2.5$, whereas in the second task, $p = 0.4$ and $R = 3$. One of these two risk tasks was randomly chosen to account for payment and earnings from the risk-elicitation task was also added to the earnings from the main task. Conducting two similar tasks with different parameters allows us to reduce measurement errors as shown in Gillen et al. (2019).

In the **Competition treatment**, all procedures were identical to the No Competition treatment, except we had two sellers competing for a single buyer. Hence, in addition to specifying her

[19]The screenshots depicting feedback that subjects received at the end of each block are presented in Section 3 of the Online Appendix.

[20]Dal Bo and Frechette (2019) use a similar method when they study infinitely repeated prisoners' dilemma games and use the strategy method.

purchasing cutoffs, the buyer needed to indicate which seller she would buy from given the messages received from each. Four different scenarios could occur: either both sellers sent message m_0, or both sellers sent message m_1, or seller 1 sent m_0 and seller 2 sent m_1, or seller 1 sent m_1 and seller 2 sent m_0. For each of these four cases, the buyer specified the probability, a number between 0 and 1, that she wants to be matched with seller 1 (with the remaining probability she was matched with seller 2). Sellers who were not matched were paid zero, whereas those who were matched received payoffs identical to those in the No Competition treatment, conditional on their specified strategy and that of the buyer. We again used the block structure for payoffs here and paid either the game payoffs or the belief payoffs for one randomly selected block.

In the **Monetary treatment**, although all the procedures were identical to the No Competition treatment, the payoffs did not reflect the psychological costs. Instead, the participants simply played the game with payoffs described in Figure 1. So, a seller was asked to specify the message that would be sent to the buyer for each possible product quality he might possess, and a buyer was asked to specify her purchasing decision for each of the two messages she could receive from the seller. We also elicited buyers' and sellers' beliefs as before. Our experimental design is summarized in Table 1.

Table 1: Experimental Design

Treatment	Number of sessions	Number of subjects
Monetary	3 sessions	58 subjects: 29 Buyers and 29 Sellers
No Competition	3 sessions	52 subjects: 26 Buyers and 26 Sellers
Competition	4 sessions	69 subjects: 23 Buyers and 46 Sellers

3.2 Discussion of experimental design choices

Inducing guilt and lying aversion in the lab. Psychological games take their name from the fact that decision-makers may be affected by their beliefs about others and their beliefs about others' beliefs about them (second-order beliefs). These beliefs can create a variety of emotions on the part of the decision-maker, which would affect how one plays the game. Hence, to properly test a psychological game in the lab, these emotions must be controlled or inferred ex-post given subject behavior.

One unique innovation of our experiment is that we induce psychological payoffs in the No Competition and the Competition treatments. Specifically, we impose costs on the seller whenever he lies to the buyer and disappoints her. We also impose the disappointment costs on the buyer when she is misled by the seller. These lying, guilt, and disappointment costs are induced and take on different values depending on the player type, in contrast to other experiments in which such costs are typically inferred.[21]

Our approach of inducing psychological costs is in line with the standard notion of induced value as originated by Smith (1976), in which an experimenter assigns payoffs to outcomes in such a way

[21] Given that we induce psychological payoffs, our focus is not on assessing whether real-world agents suffer from guilt, disappointment or lying aversion, but rather how their behavior changes in the presence of such psychological motives. By inducing them, we can observe whether behavior in the face of these motives is consistent with what our model predicts.

that any subject whose utility function is monotonic in lab payoffs will act as if they are maximizing the induced utility function. However, inducing guilt, disappointment, and lying aversion is tricky because people walk into the lab with their own homegrown attitudes toward lying and deceit, and these attitudes may be overlaid on top of or exceed the penalties we impose. This could imply a lack of control.

Our experimental design addresses these concerns and allows us to test whether inducing psychological costs works. To do this note that from a theoretical point of view, either the penalty we impose for lying, disappointment, and guilt is binding or it is not. What we mean by binding is that either the imposed penalty is more severe than the one subjects would impose on themselves given their homegrown attitudes or it is less severe. If it is more severe, we are in control of the subjects' behavior, because our penalties are sufficiently large to be the determining factor in subjects' calculations. Using the language of Smith (1976), this means the Dominance Principle is satisfied; that is, the reward medium dominantly determines changes in the subject's utility. The other case to consider is when a subject's moral aversion to guilt, disappointment, and lying is greater than the penalties we impose. Our experimental design allows us to detect this case by comparing behavior in the Monetary treatment with that in the No Competition treatment, which differs by the inclusion of psychological payoffs in the latter case. Specifically, if subjects had greater resistance to lying or misleading others than the one we imposed in the No Competition treatment, we should observe the same amount or strictly less lying in the Monetary treatment than in the No Competition treatment. The difference between lying in these two treatments would indicate the degree to which induced costs crowd out the homegrown psychological costs.

In the Results section, we address this point. We show that without induced psychological costs, sellers lie to a far greater extent than they do when such costs are induced (Table 3), providing validation that our technique increases experimental control over psychological forces central to our behavioral model.

Finally, the comparison between the No Competition and the Competition treatments remains valid regardless of whether self-imposed costs are larger or smaller than those induced in our experiments. The reason is that we use the same experimental technique in both treatments and have no reason to believe self-imposed psychological costs should respond to the number of sellers.

Eliciting beliefs and actions in psychological games. In psychological games, payoffs are a function of both actions and beliefs. Therefore, to properly test psychological games in the lab, one may need to elicit subjects' actions and their beliefs about the actions of others. However, asking subjects to report beliefs has two effects: first, it affects subjects' payoffs in the belief elicitation exercise itself, and, second, it affects their payoff in the game, since game payoffs depend on stated beliefs. The second effect implies that subjects might have a strong incentive to report zero beliefs in an effort to eliminate their guilt and disappointment costs. This will increase a subject's payoff in the game since such feelings are subtracted from their material payoffs.[22] In other words, if one takes psychological games seriously and wants to test their equilibrium in the lab, this can be a severe problem.

In our experiments, subjects are paid both for the beliefs they state and for the actions they choose. In particular, subjects are paid either for their performance in the game or for the accuracy

[22] In our setup, a seller can inflate his game payoff by reporting a zero second-order belief about message m_1 to minimize the guilt payoff. A buyer can do the same by reporting a zero first-order belief about m_1 to minimize the disappointment payoff.

of their beliefs in one randomly selected block of the experiment.[23] The beliefs are elicited using the quadratic scoring rule which penalizes them based on the difference between their stated belief for m_i and the actual observed probability that message m_i comes from a high-quality seller. Similarly, the sellers are penalized based on the difference between the belief they state and the belief buyers state for the same message m_i. We set payoffs for this quadratic scoring rule so that there is practically no incentive to report false beliefs by making the subjects indifferent at the margin between reporting false beliefs or not.[24] Theoretically, this means that they trade off their payoff in the game versus their payoff in the belief elicitation procedure. To help subjects comprehend this, we explained this to them and also told them they had no incentive to report beliefs falsely if they wanted to maximize the expected dollar payoff in the experiment. This method of announcing to the subjects facing a complicated belief elicitation procedure that truth-telling is the optimal thing to do is the most effective way to elicit true beliefs in the lab as shown in the recent influential paper by Danz et al. (2021) and is commonly done in the experiments. Given that the highest distortion in beliefs leads to a truly minimal increase in players' payoffs and the costs of contemplating a deviation from truthful reporting far outweigh any benefit from doing so, we believe that our message to subjects stating that they had no incentive to report beliefs falsely if they wanted to maximize the expected dollar payoff in the experiment was practically honest.

We now turn to our data to investigate whether subjects reported beliefs strategically or not. Remember the problem is that if subjects are strategic in reporting their beliefs they will report zero beliefs in an effort to increase their game payoffs (at the expense of their belief-elicitation payoffs). Our data suggest that this hardly ever happens. In fact, buyers never reported zero beliefs for message m_1 in any treatment, while sellers did so just a few times (less than 2%) in the Competition treatment and never in the No Competition treatment. This evidence strongly suggests that while the task of eliciting actions and beliefs in a psychological game is theoretically challenging, it was not something that dawned on our subjects nor should it since we arranged payments so as to create 'almost complete' incentive compatibility. This is, however, an issue that needs to be addressed when one tests psychological games in the lab by inducing psychological payoffs.

[23] This randomization diminishes hedging motives which lead to stating beliefs that are not compatible with one's actions (see Costa-Gomez and Weizsacker (2008)).

[24] To illustrate this approach, consider a buyer in the No Competition treatment, who believes that there is $p(m_i)$ chance that message m_i is sent by a high-quality seller and, instead, reports r in the experiment. We used the quadratic scoring rule to incentivize belief reports. As we show in the Online Appendix, this means that a buyer's expected payoff in the belief task can be written as

$$\mathbb{E}\Pi^{\text{beliefs}}\left(p(m_i), r\right) = p(m_i) \cdot \left[c_1 - 2c_2(1-r)^2\right] + (1 - p(m_i)) \cdot \left[c_1 - 2c_2 r^2\right],$$

where (c_1, c_2) are pre-specified parameters. The buyer's payoff from playing the game is

$$\mathbb{E}\Pi^{\text{game}}\left(p(m_i), r, \omega\right) = \begin{bmatrix} 10 \cdot p(m_i) + (1 - p(m_i)) \cdot (-10 \cdot \omega \cdot r) & \text{if this payoff is greater than 5} \\ 5 & \text{otherwise} \end{bmatrix}.$$

A risk-neutral buyer should report belief r^* that maximizes his overall expected payoff

$$\mathbb{E}\Pi^{\text{Buyer}}\left(p(m_i), r, \omega\right) = \frac{1}{2} \cdot \mathbb{E}\Pi^{\text{belief}}\left(p(m_i), r\right) + \frac{1}{2} \cdot \mathbb{E}\Pi^{\text{game}}\left(p(m_i), r, \omega\right).$$

The highest distortion in beliefs is equal to $\max|p(m_i) - r^*| = \frac{5}{2c_2} \cdot \left(1 - \frac{1}{\sqrt{2}}\right)$ and does not exceed 1.5% given the parameters used in the experiment, i.e., $(c_1, c_2) = (100, 50)$. Moreover, it results in a minimal increase in the buyer's payoff relative to reporting the true belief. In other words, our payment scheme is "practically" incentive compatible. We refer the reader to Section 4 in the Online Appendix, where we describe this procedure in detail.

4 Results

This section describes the performance of markets across our three treatments. We start by investigating the effects of induced psychological payoffs and competition on trade frequencies and participants' payoffs. We then document buyers' and sellers' strategies and show which psychological types of sellers are mostly affected by sellers' competition. We conclude this section by comparing the outcomes in each market with those predicted by the theory to see if any of the equilibria organize observed data in a satisfactory manner.[25]

4.1 Trade and Welfare

Figure 2 depicts the frequency of buyers' purchasing decisions in each treatment and the quality of the purchased goods. While all three treatments display similar outcomes in the first half of the experiment, once subjects have had the time to learn the game their behavior diverges and we document significant differences in the performance of markets across the three treatments.

Figure 2: Aggregate Outcomes, by treatment

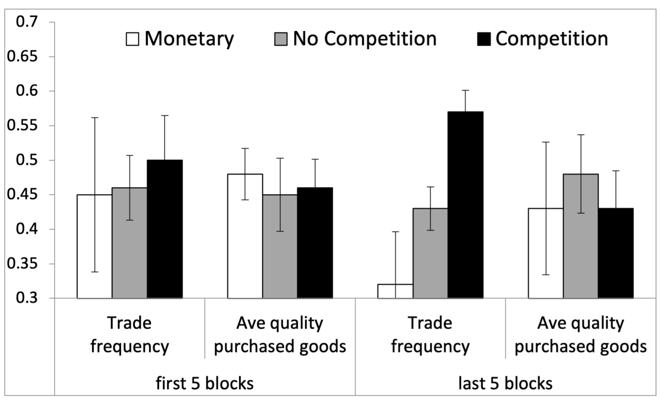

Notes: The left panel focuses on the first 5 blocks and the right panel on the last 5 blocks of the experiment. In each panel, we depict the trade frequency and the likelihood that the product was high quality conditional on the product being purchased. Bars indicate 95% confidence intervals using robust standard errors, which are computed by clustering observations by session.

[25]Throughout this section, we use regression analysis to compare average outcomes between two groups (be that two treatments or two different types of sellers). Specifically, we run random-effects GLS or LOGIT regressions (depending on the nature of the dependent variable) in which we regress the variable of interest (e.g., the purchasing decision of buyers or the quality of the sold product) on a constant and a dummy variable that indicates one of the considered groups (i.e., two treatments or two messages), while clustering observations by sessions to account for potential interdependencies of observations within a session. We say that there is a significant difference between the two considered groups if the estimated coefficient on the dummy variable is significantly different from zero, and we report the p-value associated with it. Most of the analysis focuses on the last five blocks of each experimental session because subjects often learn the game by playing it. For this reason, the data from the first iterations of the game tends to be noisier. We also present subjects' behavior in the first five blocks in several figures and tables in the main text of the paper and in the Online Appendix to highlight changes in subject behavior.

The trade frequency is significantly higher in the markets with both material and psychological payoffs compared to those with only material payoffs (Monetary vs No Competition, last 5 blocks: $p = 0.003$), while the average quality of sold goods is marginally higher in the former treatment but not significantly so (Monetary vs No Competition, last 5 blocks: $p = 0.119$). The additional introduction of competition between sellers into the market with psychological payoffs further increases the trade frequency but lowers the quality of sold goods (No Competition vs Competition, last 5 blocks: $p < 0.001$ for trade frequency and $p = 0.098$ for quality of purchased goods). In fact, the competition between sellers undoes the benefits of induced psychological payoffs in the sense that it leaves the buyers with goods that are comparable in quality to those purchased in the Monetary treatment (Monetary vs Competition, last 5 blocks: $p = 0.995$).

Table 2: Welfare of Market Participants, by treatment

	first 5 blocks		last 5 blocks	
	Buyers' Payoffs	Sellers' Payoffs	Buyers' Payoffs	Sellers' Payoffs
Monetary	4.90 (0.09)	9.79 (0.62)	4.76 (0.16)	8.65 (0.55)
No Competition	4.39 (0.15)	8.90 (0.36)	4.57 (0.11)	8.84 (0.19)
Competition	4.30 (0.13)	3.42 (0.22)	3.80 (0.15)	3.17 (0.17)
Monetary vs No Competition	$p = 0.001$	$p = 0.215$	$p = 0.342$	$p = 0.742$
Monetary vs Competition	$p < 0.001$	$p < 0.001$	$p < 0.001$	$p < 0.001$
No Competition vs Competition	$p = 0.534$	$p < 0.001$	$p < 0.001$	$p < 0.001$

Notes: The first three rows report average payoffs of buyers and sellers with robust standard errors in parentheses. Standard errors are clustered at the session level. The last three rows report the p-values comparing average payoffs across pairs of treatments as described in Footnote 25.

Next, we consider participants' payoffs. We abstract away from the payoffs that subjects earn in the belief-elicitation tasks and consider only the game payoffs. Table 2 reports the average payoffs of buyers and sellers in each treatment.

First, we compare the markets with psychological payoffs with the baseline market, in which players have only material payoffs. Recall that in the markets with psychological payoffs, the welfare of market participants is affected by two components: the market performance, i.e., the trade frequencies and the trade qualities, and the psychological costs. The second component is not present in the markets with only material payoffs. In the No Competition treatment, buyers purchase more goods of marginally higher quality, which increases their welfare relative to the Monetary treatment. However, this comes at the cost of experiencing disappointment at times when purchased goods turn out to be of low quality. The two effects exactly offset each other, resulting in comparable average buyers' payoffs in the two treatments (Monetary vs No Competition, last 5 blocks: $p = 0.342$). Similarly, the sellers tend to sell their goods more often in the No Competition than in the Monetary treatment, but some sellers experience guilt and/or lying costs when they do so, resulting in overall comparable average sellers' payoffs in the two treatments (Monetary vs No Competition, last 5 blocks: $p = 0.742$). The situation changes when we compare payoffs in the Competition and the Monetary treatments. Markets with psychological payoffs and competition feature significantly lower payoffs both for buyers and for sellers as compared with the baseline markets (Monetary vs Competition, last 5 blocks: $p < 0.001$ for both comparisons).

Second, we compare welfare in the two markets with psychological payoffs. Table 2 shows that both buyers and sellers suffer from the presence of competition (No Competition vs Competition, last 5 blocks: $p < 0.001$ in both comparisons). These losses are large in magnitude and constitute the main punchline of our paper as summarized in Result 1 below. In Section 5.4 we come back to the welfare comparisons across treatments and decompose the reduction in welfare into changes in trade outcomes driven by players' strategies and changes in players' beliefs, which determine players' psychological costs.[26]

Result 1:

(1a) *Market performance. Trade in markets without induced psychological payoffs is infrequent and when it happens most of the purchased goods are low-quality. The introduction of psychological payoffs increases both the trade and marginally increases the quality of purchased goods. The further introduction of competition between sellers into markets with psychological payoffs results in even higher trade frequency but lowers the quality of purchased goods.*

(1b) *Welfare. Buyers' and sellers' payoffs in markets with psychological payoffs and no competition are comparable to markets without psychological payoffs. However, both buyers and sellers get significantly lower payoffs in markets that feature both psychological payoffs and competition compared with any of the two other markets that feature no competition.*

4.2 Strategies of Buyers and Sellers

In this section, we look under the hood of our results and examine the strategies used by buyers and sellers. We do this by investigating the way sellers communicate with buyers, i.e., what messages sellers attach to goods of different qualities, the way buyers translate sellers' messages into purchasing decisions, and the resulting quality of goods that come with different messages.

Table 3 shows that there is little variation in sellers' communication strategies across our three treatments when subjects own a high-quality product; the vast majority of them send messages m_1.[27] It is the behavior of sellers with low-quality goods that differs. For example, the introduction of psychological payoffs reduces the lying frequency for low-quality sellers from 62% in the Monetary treatment to 24% in the No Competition treatment (last 5 blocks: $p < 0.001$). However, once the competition between sellers is introduced, the disciplining effect of psychological payoffs vanishes, and sellers with low-quality products lie as much in the Competition treatment as they do in the Monetary treatment (last 5 blocks: $p = 0.841$).

Turning to the buyers' behavior, we find that buyers' purchasing decisions after receiving message m_1 are quite similar across the treatments in the first half of the experiment.[28] However, in the Monetary treatment, the buyers learn to purchase goods less after message m_1 as they gain experience (56% in the first 5 blocks vs 39% in the last 5 blocks, $p < 0.001$). There is no such declining trend in the markets with psychological payoffs. On the contrary, buyers in the Competition treatment learn to purchase goods with an m_1 label more often in the second half compared to the first half of the experiment ($p < 0.001$). As a result, in the second half of the experiment, buyers are less likely to purchase a product after message m_1 in the Monetary treatment compared

[26]This exercise requires first documenting the changes in players' strategies and beliefs across the two treatments, which is why we perform it at the end of the paper.

[27]Focusing on the last 5 blocks of the experiment, the frequency of sending an m_1 message by a high-quality seller in any pair of treatments is not significantly different with $p = 0.971$ for Monetary vs No Competition treatment, $p = 0.095$ for Monetary vs Competition treatment, and $p = 0.401$ for No Competition vs Competition treatment.

[28]In the first half of the experiment, we detect no significant difference in the likelihood of buying a product conditional on receiving the m_1 message between any pair of treatments, $p > 0.10$ in all three pairwise comparisons.

Table 3: Messages and Purchasing Decisions, by treatment

		Monetary	No Competition	Competition
first 5 blocks				
Sellers' communication behavior	$\Pr[m_1 \mid q_H]$	0.96 (0.02)	0.90 (0.06)	0.85 (0.02)
	$\Pr[m_1 \mid q_L]$	0.60 (0.07)	0.28 (0.03)	0.59 (0.02)
Buyers' purchasing behavior	$\Pr[\text{Buy} \mid m_1]$	0.56 (0.07)	0.59 (0.03)	0.57 (0.04)
	$\Pr[\text{Buy} \mid m_0]$	0.08 (0.02)	0.32 (0.04)	0.35 (0.03)
Quality of goods for different messages	$\Pr[q_H \mid m_1]$	0.53 (0.03)	0.66 (0.03)	0.51 (0.03)
	$\Pr[q_H \mid m_0]$	0.07 (0.04)	0.07 (0.04)	0.22 (0.02)
last 5 blocks				
Sellers' communication behavior	$\Pr[m_1 \mid q_H]$	0.92 (0.01)	0.90 (0.05)	0.89 (0.01)
	$\Pr[m_1 \mid q_L]$	0.62 (0.02)	0.24 (0.03)	0.63 (0.03)
Buyers' purchasing behavior	$\Pr[\text{Buy} \mid m_1]$	0.39 (0.04)	0.56 (0.05)	0.65 (0.02)
	$\Pr[\text{Buy} \mid m_0]$	0.13 (0.06)	0.32 (0.04)	0.34 (0.03)
Quality of goods for different messages	$\Pr[q_H \mid m_1]$	0.49 (0.02)	0.71 (0.02)	0.49 (0.03)
	$\Pr[q_H \mid m_0]$	0.13 (0.02)	0.07 (0.04)	0.17 (0.02)

Notes: The average observed quantities are presented with robust standard errors in parentheses. Standard errors are clustered at the session level.

with either No Competition or Competition treatments (Monetary vs No Competition: $p = 0.001$; Monetary vs Competition: $p < 0.001$). Furthermore, the purchasing frequencies for m_1 messages are slightly higher in the Competition than in the No Competition treatment in the second half of the experiment ($p = 0.017$) and are comparable for the m_0 messages ($p = 0.627$).[29]

Finally, we translate sellers' communication strategies into the average quality of goods that come with the message m_1. Roughly 50% of these goods are actually high-quality goods in both the Monetary and the Competition treatment (last 5 blocks: $p = 0.827$). Contrary to that, in the No Competition treatment, most of such goods (71%) are actually high-quality goods (Monetary vs No Competition, last 5 blocks: $p < 0.001$; Competition vs No Competition, last 5 blocks: $p < 0.001$).

Result 2: *The sellers with low-quality goods lie the least in the markets with psychological payoffs and no competition and lie substantially more often in the other two market structures. The buyers purchase more products with m_1 labels in markets with competition and psychological payoffs than in both markets without competition.*

4.3 The Effect of Psychological Types on Strategies and Payoffs

To understand the large difference between lying frequencies of low-quality sellers in markets with psychological payoffs, we look into the sellers' communication strategies presented in Figure 3.[30]

Consistent with the averages presented in Table 3, we find that the vast majority of sellers who own a high-quality product disclose it truthfully and send message m_1. In fact, the 95% confidence

[29]In the second half of the experiment, the buyers purchase the product with the m_0 message more often in the two markets with induced psychological payoffs than in the market with only monetary payoffs (Monetary vs No

Figure 3: Communication Strategies of Sellers in Markets with Psychological Payoffs, last 5 blocks

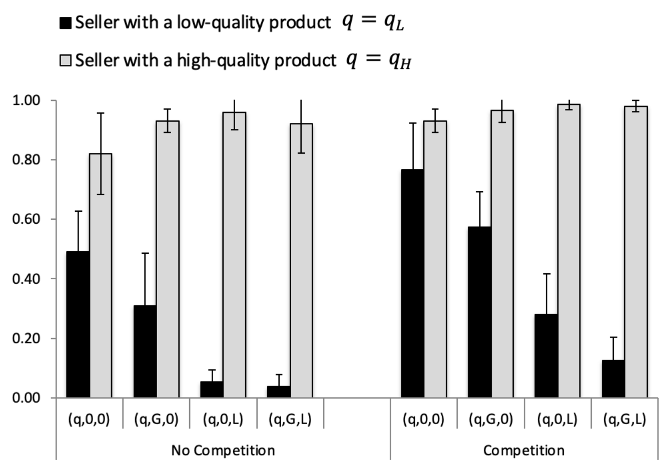

Notes: Average frequency of sending message m_1 is presented for each type of seller in each treatment in the second half of the experiment. We compute 95% confidence intervals using robust standard errors obtained by clustering observations by session.

intervals around the average frequencies of the m_1 message contain 100% for sellers with a high-quality product and a positive lie aversion in the No Competition treatment and for all sellers with a high-quality product and either guilt or lie aversion or both in the Competition treatment. The deviations of the remaining types from being 100% truthful are rather small and are not very surprising given that any small tremble would be one-sided because of the boundary.

The situation changes when we look at sellers with low-quality goods. In the No Competition treatment, we find that about 50% of low-quality sellers with no guilt and no lie aversion choose to lie, whereas other low-quality sellers lie much less. By contrast, in the Competition treatment, the sellers with types $(q_L, 0, 0)$ and $(q_L, G, 0)$ lie most of the time (more than 50%), whereas the $(q_L, 0, L)$ and (q_L, G, L) types lie much less.[31] In fact, in both treatments, we observe a monotonic decrease in the lying frequency of low-quality sellers as we move from $(q_L, 0, 0)$ to $(q_L, G, 0)$ to $(q_L, 0, L)$ to (q_L, G, L).[32]

Competition (Competition), last 5 blocks: $p = 0.047$ ($p = 0.014$).

[30] The sellers' communication strategies in the first five blocks look very similar and are presented in Figure 5 in the Online Appendix.

[31] The fact that, in the No Competition treatment, the sellers with type $(q_L, 0, 0)$ are truthful about half of the time is consistent with the idea that some participants experience an additional home-grown aversion to lying that they brought to the experiment. Interestingly, the competition between the sellers dominates this home-grown aversion to lying since the same sellers lie about 80% of the time in the markets with multiple sellers.

[32] In the last five blocks of the No Competition treatment, the $(q_L, 0, 0)$ types lie significantly more than the $(q_L, G, 0)$ types ($p = 0.004$), the $(q_L, G, 0)$ types lie significantly more than the $(q_L, 0, L)$ types ($p < 0.001$), whereas there is no significant difference between lying frequencies of $(q_L, 0, L)$ and (q_L, G, L) types ($p = 0.556$). In the last five blocks of the Competition treatment, the $(q_L, 0, 0)$ types lie significantly more than the $(q_L, G, 0)$ types ($p < 0.001$), the $(q_L, G, 0)$ types lie significantly more than the $(q_L, 0, L)$ types ($p < 0.001$), and the $(q_L, 0, L)$ types

Table 4: Which Types of Buyers and Sellers Suffer the Most from Competition in the last 5 blocks?

		No Competition	Competition	Difference
SELLERS	$(q_L, 0, 0)$	12.20 (0.84)	14.93 (1.03)	YES** ($p = 0.05$)
	$(q_L, G, 0)$	8.27 (0.54)	6.67 (0.76)	YES* ($p = 0.08$)
	$(q_L, 0, L)$	9.91 (0.69)	0.08 (1.11)	YES** ($p < 0.01$)
	(q_L, G, L)	8.07 (1.22)	-0.70 (1.94)	YES** ($p < 0.01$)
	$(q_H, 0, 0)$	7.46 (0.28)	7.83 (0.23)	NO ($p = 0.20$)
	$(q_H, G, 0)$	7.69 (0.26)	8.52 (0.41)	NO ($p = 0.15$)
	$(q_H, 0, L)$	7.90 (0.27)	7.99 (0.23)	NO ($p = 0.72$)
	(q_L, G, L)	7.56 (0.23)	7.77 (0.26)	NO ($p = 0.54$)
BUYERS	$\omega \leq 0.2$	4.15 (0.33)	4.14 (0.47)	NO ($p = 0.96$)
	$0.2 < \omega \leq 0.4$	4.11 (0.33)	3.37 (0.43)	NO ($p = 0.18$)
	$0.4 < \omega \leq 0.6$	4.74 (0.25)	4.08 (0.43)	NO ($p = 0.17$)
	$0.6 < \omega \leq 0.8$	4.78 (0.22)	2.97 (0.35)	YES** ($p < 0.01$)
	$\omega > 0.8$	5.04 (0.13)	4.54 (0.21)	YES** ($p = 0.04$)

Notes: We report the average payoffs of buyers and (selected) sellers in the last five blocks of the experiment and the robust standard error in parentheses. The last column reports the results of a statistical test comparing payoffs for a fixed type of buyer or seller in the two treatments. * and ** indicate significance at the 10% and the 5% levels, respectively.

More importantly, sellers change their behavior when competition is introduced. Sellers with low-quality goods lie significantly *more* in the Competition than in the No Competition treatment for each possible combination of their guilt and lie aversion. In the No Competition treatment, sellers with types $(q_L, 0, 0)$, $(q_L, G, 0)$, $(q_L, 0, L)$, and (q_L, G, L) sent untruthful m_1 messages 49%, 31%, 5%, and 4% of the time, respectively, while these percentages increased to 77%, 57%, 28%, and 13%, when the competition was present. Pairwise comparisons between these fractions confirm the directional results evident in Figure 3: $p = 0.011$ for types $(q_L, 0, 0)$, $p = 0.004$ for types $(q_L, G, 0)$, $p < 0.001$ for types $(q_L, 0, L)$, and $p = 0.042$ for types (q_L, G, L).

Frequent lying of sellers with high psychological costs in the Competition treatment affects their overall payoffs in the game. Table 4 addresses this point and shows which types of sellers and buyers suffer the most from the presence of competition. For the Competition treatment, we focus on the payoffs of the selected sellers since the non-selected sellers earn a fixed payoff of zero.[33]

As Table 4 shows that the selected sellers who own the high-quality product earn the same average payoffs in both treatments regardless of their psychological costs. The selected sellers with low-quality goods and at least some psychological costs are the ones who suffer from the competition. The largest losses are experienced by sellers with low-quality goods who have a strong aversion to lying and might have a strong sensitivity to guilt. Interestingly, in the markets with competition, sellers with low-quality products and a high aversion to lying who are selected by buyers earn average payoffs that are not statistically different from zero, which is the payoff of the non-selected seller. As for the buyers, those with higher disappointment aversion suffer the most losses from competition between sellers.

Result 3: *In markets with psychological payoffs, sellers with higher guilt and lying sensitivity lie less than those with lower values. Competition leads to more lying by sellers who own low-quality*

lie significantly more than the (q_L, G, L) types ($p < 0.001$).

[33]Table 1 in the Online Appendix presents the same statistics for the first 5 blocks of the experiment.

products, irrespective of their psychological costs. Furthermore, competition negatively affects the payoffs of buyers with high disappointment sensitivity and payoffs of selected sellers with low-quality goods and high lying or guilt aversion.

4.4 Equilibrium predictions

We finish this section by comparing the behavior of our market participants to the equilibrium predictions (Section 2.4). We start with the pooling equilibrium, which predicts that sellers' messages are uncorrelated with the product quality, buyers treat messages as uninformative and ignore them when making purchasing decisions, and, as a result, no trade should occur. More precisely, we define a pooling-equilibrium strategy for a seller to mean one where the seller chooses the same message regardless of the seller's type or for which there is no correlation between the message she sends and the quality of the good she is selling. A pooling-equilibrium strategy for a buyer means making the same purchasing decision irrespective of the received message in the Monetary treatment and choosing two "similar" purchasing cutoffs in the No Competition and Competition treatments, where similar means cutoffs are no more than 0.05 away from each other.

We find that the Monetary treatment is the closest among all three treatments to the pooling equilibrium described above. First, as Figure 2 indicates, trade during the last five blocks of the experiment is extremely low in the Monetary treatment (almost non-existent) and significantly less than our Competition and No Competition treatments. Second, in the last 5 blocks of the experiment, the majority of both buyers and sellers in markets without induced psychological payoffs play the pooling equilibrium strategies as defined above: 61% of the sellers and 55% of the buyers.

In contrast, the behavior of participants in the markets with psychological payoffs clearly rejects the hypothesis that they are playing the pooling equilibrium. Indeed, sellers' communication strategies are informative of the product quality (Figure 3) and buyers incorporate this information when making their purchasing decisions (Table 3). Individual level analysis confirms the same point: less than 10% of the sellers and less than 30% of the buyers in either treatment play the pooling equilibrium strategies (see Tables 2 and 3 in the Online Appendix).

Having dispensed with the pooling equilibrium as the one compatible with the No Competition and Competition data, we turn next to the two partially informative equilibria: PIE1, in which two of the four sellers' types with low-quality products lie in equilibrium, and PIE2, in which only one type does, the one that experiences no guilt and no lying aversion. We ask whether one of these two equilibria organizes data in a satisfactory manner. To do that, we present in Table 5 a subset of the PIE1 and PIE2 predictions and compare them to the observed outcomes. This subset contains those predictions that differ across these two equilibria.

The No Competition treatment conforms to some PIE2 predictions but not to all. In particular, the sellers' behavior and the average quality of goods that come with the m_1 label match theoretically predicted levels. However, buyers purchase goods significantly less than what PIE2 predicts. This leads to a lower frequency of overall trade. At the same time, the data in the Competition treatment is clearly incompatible with theoretical predictions of either PIE1 or PIE2. Hence, we conclude that the difference in behavior between the Competition and the No Competition treatment cannot be attributed to subjects selecting different equilibria in these treatments.

Result 4: *Majority of buyers and sellers play the pooling equilibrium strategies in markets with only material payoffs. The markets with psychological payoffs and no competition conform to predictions of the most informative equilibrium (PIE2) with the exception of lower-than-predicted*

Table 5: Aggregate Fit of Equilibrium Predictions, last 5 blocks

	Observed	Predicted		Statistical Tests	
		PIE1	PIE2	Observed vs PIE1	Observed vs PIE2
No Competition treatment					
Pr[Buy]	0.43 (0.02)	0.20	0.55	$p < 0.01$	$p < 0.01$
$\Pr[m_1\|q_L]$	0.24 (0.03)	0.50	0.25	$p < 0.01$	$p = 0.62$
$\Pr[\text{Buy}\|m_1]$	0.56 (0.05)	0.29	1.00	$p < 0.01$	$p < 0.01$
$\Pr[q_H\|m_1]$	0.71 (0.02)	0.57	0.73	$p < 0.01$	$p = 0.24$
Competition treatment					
Pr[Buy]	0.57 (0.02)	0.26	0.80	$p < 0.01$	$p < 0.01$
$\Pr[m_1\|q_L]$	0.63 (0.03)	0.50	0.25	$p < 0.01$	$p < 0.01$
$\Pr[\text{Buy}\|m_1]$	0.65 (0.02)	0.29	1.00	$p < 0.01$	$p < 0.01$
$\Pr[q_H\|m_1]$	0.49 (0.03)	0.57	0.73	$p = 0.01$	$p < 0.01$

Notes: The table reports observed and theoretically predicted frequencies of trade, m_1 messages, and purchasing decisions of buyers focusing on those that are different across the two partially informative equilibria. The last two columns report the p-values comparing observed frequencies with those predicted by PIE1 and PIE2.

trade frequency. The markets with psychological payoffs and competition between sellers do not match either of the equilibria.

5 Main forces driving markets to these outcomes

Our results so far have generated a number of puzzles that need to be discussed. The first is why buyers in markets with psychological costs but no competition are hesitant to buy goods after receiving an m_1 message yet the behavior of sellers in this treatment is basically consistent with the most informative equilibrium, i.e., is relatively honest. Second, why competition between sellers leads to more lying by sellers? Third, why do buyers purchase products with the label m_1 more often in markets with competition when sellers in those markets are more prone to lying?

In this section we attempt to offer some solutions to these puzzles but before we do let us pause and discuss the beliefs of our sellers and buyers in the Competition and the No Competition treatments. These beliefs will be the key ingredient in understanding subjects' behavior.

Beliefs. Recall that in the experiment, we elicited buyers' first-order beliefs of what messages mean and sellers' second-order beliefs regarding buyers' first-order beliefs. Specifically, buyers indicate the likelihood that message m_i was sent by a high-quality seller, while the sellers try to guess the likelihood stated by the buyers, i.e., buyers' interpretation of messages. Table 6 presents summary statistics of elicited beliefs and compares those to actual meanings of messages in the last 5 blocks of the experiment, while Table 4 in the Online Appendix replicates the same analysis for the first 5 blocks.

One interesting fact that stands out in Table 6 is that while buyers correctly predict the lying behavior of sellers in markets without competition, they grossly overestimate the honesty of sellers when competition is introduced. Thus, if competition leads buyers to be overly optimistic in their beliefs about sellers, it is not surprising that they purchase goods when they shouldn't. Specifically,

Table 6: Buyers' and Sellers' Beliefs, Buyers' Purchasing Cutoffs, and Actual Quality of Products for Different Messages, last 5 blocks

	observed				$z^B(m_i) = z^S(m_i)$	$z^B(m_i) = \Pr[q_H\|m_i]$	$z^S(m_i) = \Pr[q_H\|m_i]$
	$\omega'(m_i)$	$z^B(m_i)$	$z^S(m_i)$	$\Pr[q_H\|m_i]$			
No Competition							
message m_0	0.29 (0.03)	0.26 (0.06)	0.20 (0.04)	0.07 (0.04)	$p = 0.009$	$p = 0.001$	$p = 0.001$
message m_1	0.59 (0.05)	0.76 (0.03)	0.73 (0.02)	0.71 (0.02)	$p = 0.136$	$p = 0.068$	$p = 0.247$
Competition							
message m_0	0.31 (0.03)	0.22 (0.03)	0.25 (0.02)	0.17 (0.02)	$p = 0.318$	$p = 0.167$	$p < 0.001$
message m_1	0.62 (0.04)	0.77 (0.03)	0.70 (0.03)	0.49 (0.03)	$p = 0.004$	$p < 0.001$	$p < 0.001$

Notes: The first column records average cutoffs reported by buyers for each message, $\omega'(m_i)$, which is the highest disappointment sensitivity for which a buyer is willing to purchase the product that comes with message m_i. The second and third columns, $z^B(m_i)$ and $z^S(m_i)$, are buyers' and sellers' beliefs for message m_i. The fourth column, $\Pr[q_H|m_i]$, is the likelihood that message m_i comes from the high-quality seller estimated using the actual realizations observed in each round of each block. In all cells, the robust standard errors are reported in parentheses. The last three columns report results of statistical tests comparing buyers' and sellers' beliefs (fifth column), buyers' beliefs and the average actual frequency of high-quality sellers for different messages (sixth column), and sellers' beliefs and the average actual frequency of high-quality sellers for different messages (seventh column).

in markets without competition, the average buyers' beliefs for the m_1 message are 0.76 which are not statistically different from the veracity of such messages which is 0.71 ($p > 0.05$). Moreover, average buyers' beliefs match those predicted by the PIE2, which organizes the data in the No Competition treatment quite well.[34] However, in markets with competition between sellers, buyers do quite poorly at predicting the average quality of the product that comes with label m_1: average buyers' beliefs are almost 30 percentage points above the actual number (0.77 versus 0.49).

While sellers are never supposed to lie when sending the m_0 message, we find that in the experiment such beliefs are 0.26 and 0.22 in the No Competition and Competition treatment, respectively. Given that zero is a corner solution, unsurprisingly, we find that all the deviations are positive and move average buyers' beliefs away from the prediction of zero. Perhaps a more informative statistic might be the fraction of times that reported beliefs were close to zero, allowing for some small noise. In the last five blocks of the No Competition treatment, the majority (63% in the No Competition and 77% in the Competition treatment) of reported beliefs upon observing message m_0 are at most 5 percentage points away from zero.

Despite the difference in seller behavior across markets with and without competition, it is interesting that buyers' beliefs are remarkably similar which suggests that buyers did not adjust their beliefs to accommodate the different environment they were in. Indeed, after observing an m_0 message, the mean belief of buyers was 0.26 in the No Competition treatment and 0.22 in the Competition treatment. After receiving the m_1 message, these beliefs were 0.76 and 0.77, respectively. These beliefs are statistically indistinguishable across the two treatments ($p = 0.462$ for message m_0 and $p = 0.927$ for message m_1).[35]

[34] According to PIE2, 73% of m_1 messages come from a high-quality seller. We cannot reject the null that the observed buyers' beliefs are significantly different from this number with $p = 0.259$.

[35] The similarity in buyers' beliefs in the two treatments is consistent with the similarity in buyers' purchasing decisions. Conditional on receiving an m_0 message, buyers choose cutoffs of 0.29 and 0.31 for the No Competition

The performance of our markets also depends on the beliefs of sellers. We observe that sellers correctly predict buyers' beliefs for the m_1 message in the No Competition treatment ($p = 0.136$) and for the message m_0 in the Competition treatment ($p = 0.318$).[36] For the remaining two cases, message m_0 in the No Competition treatment and message m_1 in the Competition treatment, sellers think that buyers' beliefs are slightly lower than they actually are, but the differences are quite small in magnitude. Combining data for the No Competition treatment presented in Tables 5 and 6 we find that sellers' overall frequency of lying when they own a low-quality product is consistent with the beliefs they hold regarding buyers' interpretation of the m_1 messages.

With this background let us proceed to discuss our three puzzles.

5.1 Puzzle 1: Buyers' Beliefs and Purchasing Decisions in Markets without Competition

Why do buyers in markets with no competition hesitate to buy goods after receiving the m_1 message, especially when those messages are basically honest and buyers believe them to be so?[37]

A simple explanation can be found if we consider buyers to be risk-averse. The optimal purchasing cutoff for a risk-averse buyer is lower than that of a risk-neutral buyer with the same beliefs. In other words, risk aversion leads buyers to be more cautious in their purchasing behavior.

To test whether our data support this explanation, we use an additional measure of risk aversion (an investment task), which we collected at the end of each experimental session. In the investment task, subjects were asked to allocate a budget of 200 points between a risk-free asset that paid one point for every point invested and a risky asset that paid 2.5 or 3 points with a probability of 0.50 or 0.40 for each point invested in the investment task 1 and 2, respectively. Thus, subjects who invest the full amount in the risky asset are either risk-neutral or risk-loving, whereas lower than full investment indicates that a subject is risk-averse. In addition, we can rank subjects in terms of their risk attitudes: the lower the amount invested in the risky asset, the more risk averse she is. Our data indicate a significant correlation between buyers' purchasing cutoffs upon observing m_1 and their risk attitudes: buyers who are more risk averse set lower cutoffs for m_1 (No Competition treatment, last 5 blocks: $p = 0.026$).[38,39]

and Competition treatment in the last 5 blocks, which are not statistically different from each other ($p = 0.627$). Upon receiving the m_1 message, buyers' cutoffs in the last 5 blocks were 0.59 and 0.62 for the No Competition and Competition treatments, which are also not different ($p = 0.558$).

[36] The fact that in the No Competition treatment sellers' beliefs for m_1 match average buyers' beliefs and also match beliefs predicted by the PIE2 adds one more piece of evidence that the most informative equilibrium, PIE2, predicts the average behavior in the markets without competition quite well.

[37] The second row in Table 6 indicates that the average cutoff for m_1 message is 0.59. This is consistent with the numbers reported in Table 3 that shows that given actual realizations of types in the experiment, buyers purchased 56% of goods that came with message m_1. In contrast, the PIE2 predicts that buyers should always purchase the product that comes with the message m_1.

[38] To reach this conclusion, we use the ORIV (Obviously Related Instrumental Variables) technique developed by Gillen et al. (2019), which corrects for the measurement errors in the elicitation of risk attitudes. ORIV is an improved version of the traditional instrumental variables approach to errors-in-variables, which produces consistent coefficients, correlations, and standard errors and an estimator that is more efficient than standard instrumental variable techniques. We used the average m_1 cutoffs stated by the buyers in the last 5 blocks of the No Competition treatment to perform the ORIV estimation.

[39] Additionally, buyers' purchasing cutoffs seem to be well-calibrated with the estimated level of risk aversion. This exercise is conducted for buyers who report interior investments in the risk-elicitation tasks, i.e., for whom we can estimate the degree of risk-aversion using the standard CRRA utility function, $u(x) = \frac{x^{1-\rho}}{1-\rho}$. The average estimated coefficient of risk-aversion is equal to $\bar{\rho} = 0.236$. Focusing on the last 5 blocks of the experiment, we find that these

5.2 Puzzle 2: Seller Behavior in Markets with Competition

Our second puzzle is why do sellers lie more often in the markets with competition than in the markets without competition given that buyers behave the same way in the two markets and sellers are well aware of it?

As documented in Result 4, our participants in the Competition treatment display behavior that does not adhere to any of the equilibria predicted by the theory. Therefore, in this section, we explore the non-equilibrium behavioral forces that push sellers to lie more often in markets with multiple sellers.

Note that in markets with multiple sellers, the optimal actions of sellers depend not only on sellers' beliefs about buyers' behavior but also on their beliefs about the other sellers' actions. Such considerations are not present in markets with one seller and introduce an additional layer of strategic uncertainty.[40]

We hypothesize that sellers lie more often in markets with competition because they feel they have to do so to sell their goods. This is exacerbated by the feedback sellers get in these markets. Indeed, sellers observe all messages sent by their competitors but they only observe the quality of competitors' goods when the purchase happens. Thus, a seller does not always know whether the other seller sent the m_1 message because he truly owns a high-quality product or he is lying. What sellers do know from experience is that buyers tend to select a seller with the m_1 message when facing two different messages, and, as a result, the excluded seller gets a zero payoff. In our data, when faced with two identical messages, buyers are equally likely to select either one of the two sellers; the probability of selecting the first seller is not significantly different from 50% (last 5 blocks: $p = 0.53$). However, when a buyer receives two different messages, then she selects the seller with m_1 message in 84% cases in the last 5 blocks; this frequency is significantly different from 50% ($p < 0.001$).

Table 7 shows evidence consistent with the logic above. For low-quality sellers with different psychological costs, we regress an indicator that such a seller lies to the buyer and sends the m_1 message in a particular block of the experiment. The right-hand side variables include the behavior of the competitor seller in the previous block and the behavior of the buyer. We control for the seller's own tendency to lie for the same psychological costs in the previous block, the sellers' second-order beliefs, the buyer's purchasing frequency conditional on selecting the seller with the m_1 message, and the block number. The regression analysis shows that sellers react to competitive pressure by lying more often when their competitors send more m_1 messages irrespective of their psychological type and this effect is large in magnitude (the first row). In addition, the low-quality sellers with high psychological costs, i.e., $(q_L, 0, L)$ and (q_L, G, L), lie more often in response to observing that the buyer is more likely to engage in the interaction with the seller who reports m_1 over m_0 (the third row).

As Table 7 shows, the sellers with type $(q_L, G, 0)$, i.e., the type who suffers only from guilt aversion, react the most to the actions of their competitors. Such behavior is actually reasonable if the seller S_j with type $(q_L, G, 0)$ is unsure how his competitor S_k with the same type behaves.

buyers hold beliefs that are indistinguishable from those predicted by PIE2 for message m_1, which are predicted to be 73%. Given these beliefs and the level of buyers' risk aversion, the optimal purchasing cutoff after receiving message m_1 should be $\bar{\omega}(m_1) = 0.59$, which is very close to the purchasing cutoff of 0.62 which is the average observed purchasing cutoff for the buyers under consideration. We thank the referee for suggesting this exercise.

[40] If subjects were playing strategies that constitute an equilibrium, such strategic uncertainty would not be present, as actions of all players would be optimal given how others behave. However, this is not the case in the Competition treatment, since subjects are not playing the equilibrium, which is why the presence of strategic uncertainty is reasonable.

Table 7: Learning by Observing Actions of Other Sellers in the Competition treatment

	Dependent Variable: Indicator that seller S_j with type (q_L, g^j, l^j) reports m_1 in block t			
	$(q_L, 0, 0)$	$(q_L, G, 0)$	$(q_L, 0, L)$	(q_L, G, L)
Behavior of seller S_k in block $t-1$				
$\Pr[m^k = m_1]$	0.62** (0.09)	0.70** (0.11)	0.61** (0.09)	0.49** (0.07)
$\Pr[q^k = q_L \mid m^k = m_1, \text{Buy}]$	0.02 (0.67)	0.03 (0.06)	-0.06 (0.06)	-0.06 (0.04)
Behavior of the buyer in block $t-1$				
$\Pr[m^{S^{\text{win}}} = m_1 \mid (m_0, m_1)]$	-0.02 (0.04)	0.06 (0.05)	0.14** (0.05)	0.10** (0.04)
Nb obs	414	414	414	414
Nb subjects	46	46	46	46
Overall R-square	0.2437	0.2053	0.1883	0.2431

Notes: Random-effects GLS regressions with fixed effects for sessions and with robust standard errors clustered at the individual level. We control for the block number, for sellers' second-order beliefs in the current block, for sellers' own behavior given the same type in the previous block, and the buyer's purchasing frequency conditional on selecting the seller with the m_1 message in the previous block. ** indicates significance at the 5% level.

To see this point, consider seller S_j, who believes that seller S_k always sends the message m_1 if he owns a high-quality product and if he owns a low-quality product and has no psychological costs. The seller S_j also believes that seller S_k sends the message m_0 when he owns a low-quality product and is highly lie-averse, i.e., $l^k = L$. However, seller S_j is uncertain about the strategy played by seller S_k with type $(q_L, G, 0)$ and assigns a probability x to seller S_k lying in this situation. The question is then what the optimal response of seller S_j is. A simple calculation shows that seller S_j with type $(q_L, G, 0)$ would prefer to send the m_1 message even if x is small.[41] In other words, the presence of strategic uncertainty, which is surely a reality in both the actual markets with multiple sellers and in our experimental markets, facilitates sellers to lie more often than they believe their competitor does.

5.3 Puzzle 3: Buyer Behavior in Markets with Competition

Our final puzzle is why, if sellers are lying more in markets with competition, the buyers in those markets do not adjust their beliefs accordingly. This is puzzling since in markets without competition we find that buyers do learn how to interpret messages correctly.

We believe there are two main explanations for this. First, at the start of the experiment, subjects in the treatment with competition had a significantly higher belief in sellers' honesty compared to the treatment without competition. In the first two blocks, 54% of buyers in the Competition treatment held the belief that $z^B(m_1)$ is at least 80%, while in the No Competition treatment,

[41] The seller S_j with type $(q_L, G, 0)$ facing competition from the seller S_k prefers to send message m_1 over message m_0 as long as

$$\left(\frac{1}{2}\Pr[m^k = m_1] + \Pr[m^k = m_0]\right)\cdot\Big[\Pr[\text{Buy}|m_1]\cdot\Big(21 - G\cdot 10z^{S_j}(m_1)\cdot\mathbb{E}[\omega|m_1,\text{Buy}]\Big)+(1-\Pr[\text{Buy}|m_1])\cdot 5\Big] \geq \frac{1}{2}\Pr[m^k = m_0]\cdot 5,$$

where $\Pr[\text{Buy}|m_1] = \bar{\omega}(m_1) = 0.62$, $z^{S_j}(m_1) = 0.70$, $\mathbb{E}[\omega|m_1,\text{Buy}] = \frac{\bar{\omega}(m_1)}{2} = 0.31$, and $\Pr[m^k = m_1] = 1 - p + p(\frac{1}{4} + \frac{x}{4}) = 0.55 + 0.15x$ (see Table 6). The inequality above is satisfied for all $x \geq 0$.

only 29% shared this belief.[42] Hence, subjects initially enter the experiment with the perception that sellers are more truthful and trustworthy when competition exists, and they need to be convinced otherwise. One possible reason for this is the common belief that competition is beneficial in markets, aiding in the elimination of issues like price gouging and corruption. Regardless of its accuracy, subjects in the Competition treatment exhibit higher trust in sellers right from the start of the experiment.

These inflated beliefs might not be that harmful to buyers if they would quickly learn to adjust their interpretation of m_1 messages. Unfortunately for buyers, that is not the case in the markets with competition. Let us turn to our data for confirmation.

Table 8 presents regressions depicting how buyers adjust their beliefs regarding the m_1 message conditional on the feedback they observe in the experiment. For each treatment, the dependent variable is buyers' beliefs regarding m_1 in a particular block. The right-hand side variables include buyers' initial beliefs about m_1, their previous block beliefs about m_1, and the likelihood of purchasing a high-quality good for different messages in the previous block.

Table 8: The Evolution of Buyers' Beliefs

	Beliefs in block t	
	No Competition	Competition
Beliefs		
$z^B(m_1)$ in block 1	0.05 (0.06)	0.35** (0.07)
$z^B(m_1)$ in block $t-1$	0.57** (0.06)	0.27** (0.07)
Feedback observed in block $t-1$		
$\Pr[q_H\|m_1]$	0.07** (0.04)	0.09** (0.04)
$\Pr[q_H\|m_0]$	0.02 (0.07)	0.04 (0.04)
Nb obs	231	207
Nb participants	26	23
Overall R-sq	0.4343	0.3956

Notes: Random-effects GLS regressions with robust standard errors clustered at the individual level. $\Pr[q_H|m_i]$ is the fraction of high-quality products that came with the m_i. To account for inter-dependencies of observations that come from the same session, we include session-fixed effects. In both regressions, we control for the block number, risk attitude measures, and buyers' beliefs about m_0 in the previous block.

Table 8 shows that in both treatments interpretations of m_1 messages depend on how often buyers observe high-quality products with the m_1 label in the previous block. This experience accumulates over time and affects how much buyers trust the future m_1 messages. However, in markets with competition, beliefs are more rigid and reflect to a large extent the very first belief that buyers formed before observing any sellers' behavior. This can be seen by the positive, significant, and large in magnitude estimated coefficient on $z^B(m_1)$ in the first block in markets with competition (second column), while the same coefficient is not significant in markets without competition (first column).

The buyers' interpretation of messages affects their purchasing decisions, and this is where the sluggishness of beliefs plays an important role. Table 9 illustrates the connection between beliefs and actions. To do that, we examine individual purchasing cutoffs for an m_1 message averaged over the last 5 blocks of the experiment, and ask how these average cutoffs depend on changes in beliefs about the content of an m_1 message between the two halves of the experiment, the difference in

[42] Figure 6 in the Online Appendix presents the cumulative distribution functions of these initial beliefs.

interpretations of the m_1 and m_0 messages, and the probability of choosing a seller with the m_1 message over the m_0 message when multiple sellers are present in the market.

Table 9: Buyers' Purchasing Decisions

Dependent Variable: Purchasing cutoff for an m_1, last 5 blocks		
	No Competition	Competition
$z^B(m_1)\vert_{\text{last 5}} - z^B(m_1)\vert_{\text{first 5}}$	0.66** (0.13)	1.23* (0.44)
$z^B(m_1)\vert_{\text{first 5}} - z^B(m_0)\vert_{\text{first 5}}$	0.27 (0.27)	0.56* (0.22)
$\Pr[m^{S^{\text{win}}} = m_1 \vert (m_0, m_1)]_{\text{last 5}}$		0.92** (0.09)
Nb obs	26	23
R-sq	0.3822	0.6528

Notes: Linear regressions with robust standard errors clustered at the session level. The right-hand side variable is the difference in average beliefs for m_1 message between the last 5 and the first 5 blocks of the experiment. The second one is the difference in average beliefs between m_1 and m_0 messages in the first 5 blocks. The last right-hand side variable is the chance that the buyer in the Competition treatment chooses a seller with the m_1 message over the m_0 message if two messages are different, averaged over the last 5 blocks. We control for the difference in beliefs between the two messages in the last 5 blocks of the experiment and the risk attitude of buyers.

In the markets without competition (first column), buyers' purchasing decisions are strongly affected by the change in beliefs regarding an m_1 message, which reflects their actual experiences in the markets and not just the initial beliefs. At the same time, in the markets with competition (second column), both actual experiences and initial beliefs play a role but these links are weaker as can be seen by a drop in the significance of both variables; both effects are marginal at the 10% level. The main determinant of buyers' purchasing cutoffs in these markets is how likely a buyer is to choose a seller with an m_1 message over the one with an m_0. In a sense, this last variable captures a buyer's level of skepticism towards an m_1 message relative to an m_0 one. Our data shows that this level of skepticism is positively and strongly correlated with buyers' beliefs about the informational content of messages.[43]

The slow belief updating in the Competition treatment may not be irrational, given the feedback buyers receive. Consider a buyer who receives two m_1 messages and randomly chooses one seller to make a purchase. Suppose the purchased product turns out to be of low quality. How should this buyer update her belief about the truthfulness of an m_1 message in general? The posterior belief, in this case, may not change too much because the other seller, whose offer was rejected, may have been truthful. The buyer gets no feedback about the goods she did not purchase, and, as a result, there is no reason to revise beliefs about the truthfulness of the other seller. Hence, the reduction in the trustworthiness of m_1 messages in a competitive market may not be as severe as it is in a non-competitive market, where a single deceptive message serves as a clear signal of the trustworthiness of m_1 messages. We formalize this argument in Section 6 of the Online Appendix.

[43] In particular, for each buyer in the Competition treatment, we measure the average likelihood of choosing a seller with an m_1 message over an m_0 one separately for the first 5 and the last 5 blocks of the experiment. We then compute the correlation between this measure and the difference in beliefs associated with an m_1 and an m_0 messages, separately for the first 5 and the last 5 blocks. The correlation is 0.57 ($p = 0.004$) for the first 5 blocks and 0.41 ($p = 0.053$) for the last 5 blocks.

5.4 Putting Things Together

We finish this section by looking at the informativeness of messages in the presence and absence of competition. As defined in Section 2.4, informativeness of messages is the difference in beliefs upon observing an m_1 and an m_0 message, and it tells us how much more likely message m_1 is to be sent by a seller with a high-quality product quality than message m_0. Figure 4 presents the *perceived informativeness*, which captures buyers' beliefs, and the *actual informativeness*, which depicts the correct market interpretation of messages derived from sellers' behavior.

Figure 4: Perceived and Actual Informativeness of Messages

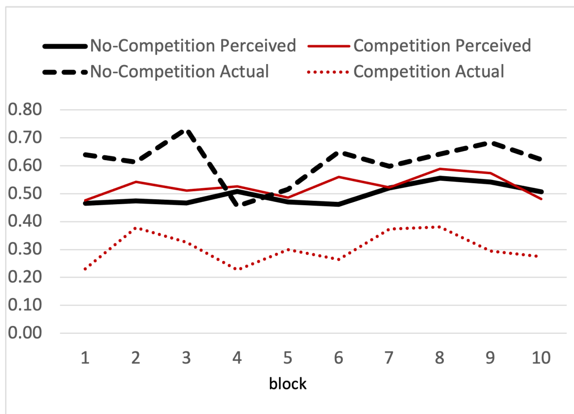

Notes: The solid lines depict the difference between buyers' beliefs about message m_1 and message m_0 averaged in a block in each treatment. The dotted lines are the difference between probabilities that the product is high quality conditional on message m_1 versus m_0 for actual realized trades.

The comparison between treatments is quite stark and corroborates the mechanism identified above. In the game without competition, buyers' interpretation of messages is close to the actual meanings of messages. Although a statistically significant difference exists between perceived and actual informativeness in a few initial blocks of the experiment, this difference by and large disappears in the later blocks.[44] By contrast, in the game with competition, the gap between what buyers think messages mean and what they actually mean is large. This gap is persistent and is not mitigated by learning.[45] In other words, competition among sellers diminishes the informational content of messages, but buyers do not realize this is happening and continue to trust messages more than they should.

Bringing all pieces of the story together, we find that markets without competition feature correct average buyer beliefs from the start and higher buyer responsiveness to sellers' strategies

[44] We compare the distribution of perceived informativeness in the No Competition treatment, estimated for each buyer in a block, with the average actual informativeness in the same block and obtain the following $p-$values for blocks 1 - 10: $p = 0.014$, $p = 0.005$, $p < 0.001$, $p = 0.374$, $p = 0.423$, $p = 0.008$, $p = 0.229$, $p = 0.124$, $p = 0.038$, $p = 0.092$.

[45] We perform the same statistical analysis for the Competition treatment as the one reported in Footnote 44 and obtain the following $p-$values for each block: $p < 0.001$, $p = 0.073$, $p = 0.150$, $p = 0.004$, $p = 0.107$, $p < 0.001$, $p = 0.025$, $p < 0.001$, $p = 0.002$, $p = 0.014$.

via the feedback they observe. In these markets, beliefs are the main driving force of buyers' purchasing decisions, and even buyers who hold initially wrong beliefs adjust them by responding to what transpires in the market. On the contrary, buyers in the markets with competition mistakenly believe from the start of the experiment that the m_1 messages are associated with a high average product quality and the follow-up experience does not convince them otherwise. In general, both beliefs and purchasing cutoffs are more rigid and less responsive to experience in markets with competition, which is why buyers never learn the true meaning of an m_1 message when multiple sellers compete with each other.

Welfare Decomposition. As we saw in Section 4.1, both the buyers and sellers suffer from the presence of competition in the markets with psychological payoffs. We now ask to what extent the reduction in welfare is driven by players' strategies, which determine the quality and the frequency of trade, versus players' beliefs, which determine psychological costs.[46]

Competition reduces buyers' welfare in markets with psychological payoffs for two primary reasons. First, these markets feature more inefficient trade, and second, buyers experience higher disappointment costs due to miscalibrated beliefs compared to markets without competition.

To assess the relative importance of these two channels, we perform the following calculation. For the markets with competition, we compute buyers' payoffs in the hypothetical scenario in which buyers' beliefs match the empirical interpretation of messages observed in markets with competition while holding all other players' strategies fixed. This calculation shows that the miscalibrated beliefs in markets with competition are accountable for only 13% of the decrease in buyers' welfare compared to markets without competition, where beliefs are, on average, accurate. This result implies that inefficient trade is the main source of the reduction in buyers' welfare in markets with competition.

In a similar spirit, we decompose the reduction in sellers' welfare into three parts. First, each individual seller is less successful at trade when he has to compete with another seller since the buyer selects only one seller to deal with. Second, sellers lie more often in markets with competition, and, as a result, some seller types pay higher lying costs. Third, sellers correctly anticipate buyers' beliefs regarding messages' interpretations, but buyers' interpretations are generally wrong in the markets with competition, which means some sellers suffer higher guilt costs relative to what they would have suffered if buyers held correct beliefs.

To assess the relative importance of these three channels, we hold fixed players' strategies and sellers' beliefs as they were in the markets with competition and compute what would be the sellers' payoffs in these markets if they faced no competition from another seller. This calculation reveals that the first channel, the competition between sellers, is the main force, which accounts for more than 90% of the reduction in sellers' welfare in markets with competition.

Feedback in markets with and without competition. We finish by noting that the feedback we provide in our experiment is comparable to what real markets have to offer. In this sense, our results about the sluggishness of buyers' beliefs in markets with competition and the competitive pressure among sellers to lie when they face competition are likely to be present in markets outside the laboratory setting. Indeed, the market structure affects how much buyers and sellers can learn from the feedback they observe. In both markets with and without competition, the sellers can learn the buyers' behavior pretty well. This is consistent with our observation that the sellers' beliefs

[46]We thank the referee for suggesting this exercise and provide all calculations for this exercise in Section 7 in the Online Appendix.

about buyers' interpretation of messages are on average correct in both markets. Moreover, in markets without competition, the buyers learn everything there is to learn about the sellers except for their psychological types, which allows them to calibrate their beliefs and anticipate the average quality of products with different labels. The situation is different in markets with competition because buyers only learn things about the selected seller and have a limited understanding of the strategy used by the seller they chose not to interact with. Our data shows that this results in buyers' inability to correctly anticipate what messages mean and a persistence in believing that competition between sellers is beneficial for them.

6 Conclusions

In this paper, we study the impact of introducing both psychological payoffs and competition into a communication (market) game and investigate their consequences for market outcomes and welfare. Specifically, we look at sellers who suffer a cost when they lie and/or mislead buyers into purchasing subpar goods, and buyers who suffer from disappointment whenever they are tricked into buying such goods.

In contrast to previous experimental work on psychological games, we impose on our participants the costs of lying, guilt, and disappointment since our focus is on testing the equilibrium model which takes these costs into account. Doing so allows us to control these psychological payoffs experimentally to some extent and evaluate their comparative static effects.[47]

We find that markets with psychological costs feature more trade of goods with marginally higher quality compared to the markets with material payoffs only. The introduction of competition between sellers in markets with psychological payoffs, however, undoes these benefits and leads to lower welfare for both buyers and sellers. While in such markets more goods are sold, considerably more of these goods are of low quality.

We analyze the mechanisms underlying these aggregate results and identify the main behavioral forces preventing competition from curbing the lying in these markets. We find that competitive pressures, especially in a winner-take-all situation such as ours, encourage sellers to misrepresent and lie more, even if they suffer from the psychological costs of doing so. The sellers' propensity to lie more is reinforced by the behavior of our buyers who fail to understand the lies the sellers tell them in such a competitive environment. Our buyers seem to genuinely believe competition is beneficial for their welfare and fail to change the way they interpret messages in markets with competition. The abundant feedback and experience we provide our subjects do not correct for the misperception of messages our buyers demonstrate in the game with competition. Consequently, sellers take advantage of such blind faith on the buyers' part and peddle lower-quality products indiscriminately.

One natural question that arises is what type of intervention may remedy the deleterious impact of competition in these markets. An obvious one is sellers' reputations which were not a feature of our experiment but do exist in the real world. Such reputations are developed in settings where

[47] As we have shown, our experimental design was successful in manipulating subjects' payoffs even if one assumes they arrive in the lab with their own aversion to lying and guilt. This is evident from the fact that in our treatments with psychological costs, we observe what appear to be partially informative equilibria that exist only when sellers experience psychological costs. At the same time, in the treatment in which we eliminate psychological costs and induce only the material ones, we do not see partially informative equilibria being played. In other words, the presence of psychological costs is beneficial in markets without competition, because they facilitate trade, which does not occur in the pooling equilibria of markets with only monetary payoffs.

interactions are repeated and where buyers can observe the outcomes of other buyers and their experiences. This should help buyers learn the product quality of sellers they have not bought from and hence add valuable information. Such information may be available on the web via customer reviews which may be of great assistance in steering these markets in the right direction and policing the behavior of those sellers who are less burdened by lying and guilt aversion.[48]

References

Abeler, J., D. Nosenzo, and C. Raymond (2019): "Preferences for Truth-Telling," *Econometrica*, 87(4), 1115–1153.

Battigalli, P. and M. Dufwenberg (2007): "Guilt in Games," *American Economic Review*, 97(2), 170–176.

——— (2009): "Dynamic psychological games," *Journal of Economic Theory*, 144, 1–35.

——— (2022): "Belief-Dependent Motivations and Psychological Game Theory," *Journal of Economic Literature*, 60(3), 833–882.

Born, A. (2020): "Promise Competition," .

Cabral, l. and A. Hortacsu (2010): "The Dynamics of Seller Reputation: Evidence from eBay," *Journal of Industrial Economics*, 58(1), 54–78.

Casella, A., N. Kartik, L. Sanchez, and S. Turban (2018): "Communication in Context: Interpreting Promises in an Experiment on Competition and Trust," *Proceedings of the National Academy of Sciences*, 115(5), 933–938.

Charness, G. and M. Dufwenberg (2006): "Promises and Partnership," *Econometrica*, 74, 1579–1601.

——— (2011): "Participation," *American Economic Review*, 101, 1213–1239.

Chen, Y., N. Kartik, and J. Sobel (2008): "Selecting Cheap-Talk Equilibria," *Econometrica*, 76(1), 117–36.

Costa-Gomez, M. and G. Weizsacker (2008): "Stated Beliefs and Play in Normal-Form Games," *Review of Economic Studies*, 75(3), 729–762.

Crawford, V. and J. Sobel (1982): "Strategic Information Transmission," *Econometrica*, 50(6)), 1431–1451.

Dal Bo, P. and G. Frechette (2019): "Strategy Choice in the Infinitely Repeated Prisoner's Dilemma," *American Economic Review*, 109(11), 3929–3952.

Danz, D., L. Vesterlund, and A. Wilson (2021): "Belief Elicitation and Behavioral Incentive Compatibility," *American Economic Review*, 112(9), 2851–2883.

Erat, S. and U. Gneezy (2012): "White Lies," *Management Science*, 58(4), 723–733.

Fischbacher, U. (2007): "z-Tree: Zurich Toolbox for Ready-made Economic Experiments," *Experimental Economics*, 10(2), 171–178.

Fischbacher, U., C. Fong, and E. Fehr (2009): "Fairness, errors and the power of competition," *Journal of Economic Behavior Organization*, 72, 527–545.

Geanakoplos, J., D. Pearce, and E. Stacchetti (1989): "Psychological games and sequential rationality," *Games and Economic Behavior*, 1, 60–79.

Gillen, B., E. Snowberg, and L. Yariv (2019): "Experimenting with Measurement Error: Techniques with Applications to the Caltech Cohort Study," *Journal of Political Economy*, 127(4), 1826–1863.

Gneezy, U. (2005): "Deception: The Role of Consequences," *American Economic Review*, 95(1), 384–394.

[48] Cabral and Hortacsu (2010) found that eBay sellers' reputations are instrumental in weeding out bad actors.

Gneezy, U. and J. Potters (1997): "An Experiment on Risk Taking and Evaluation Periods," *The Quarterly Journal of Economics*, 112(2), 631–645.

Goeree, J. and J. Zhang (2014): "Communication Competition," *Experimental Economics*, 17(3), 421–438.

Huck, S., G. Lunser, and J.-R. Tyran (2012): "Competition fosters trust," *Games and Economic Behavior*, 76, 195–209.

Hurkens, S. and N. Kartik (2009): "Would I lie to you? On social preferences and lying aversion," *Experimental Economics*, 12, 180–192.

Jin, G., M. Luca, and D. Martin (2021): "Is No News (Perceived As) Bad News? An Experimental Investigation of Information Disclosure," *American Economic Journal: Microeconomics*, 13(2), 141–173.

Kartik, N. (2009): "Strategic Communication with Lying Costs," *The Review of Economic Studies*, 76(4), 1359–1395.

Keck, S. and N. Karelaia (2012): "Does competition foster trust?" *Experimental Economics*, 15, 204–228.

Nyarko, Y. and A. Schotter (2002): "An Experimental Study of Belief Learning using Elicited Beliefs," *Econometrica*, 70(3), 971–1005.

Sanchez-Pages, S. and M. Vorsatz (2007): "Enjoy The Silence: An Experiment on Truth-Telling," *Experimental Economics*, 12(2).

Schotter, A., A. Weiss, and I. Zapater (1996): "Fairness and survival in ultimatum and dictatorship games," *Journal of Economic Behavior and Organization*, 31(1), 37–56.

Smith, V. (1976): "Experimental Economics: Induced Value Theory," *American Economic Review*, 66(2), 274–279.

Sobel, J. (2020): "Lying and Deception in Games," *Journal of Political Economy*, 128(3), 907–947.

Vanberg, C. (2008): "Why Do People Keep Their Promises? an Experimental Test of Two Explanations," *Econometrica*, 76(6), 1467–1480.

Vespa, E. and A. Wilson (2016): "Communication with multiple senders: An experiment," *Quantitative Economics*, 7, 1–36.